

Statistical Query Lower Bounds for Tensor PCA

Rishabh Dudeja

*Department of Statistics
Columbia University
New York, NY 10027, USA*

RD2714@COLUMBIA.EDU

Daniel Hsu

*Department of Computer Science
Columbia University
New York, NY 10027, USA*

DJHSU@CS.COLUMBIA.EDU

Editor: Gabor Lugosi

Abstract

In the Tensor PCA problem introduced by Richard and Montanari (2014), one is given a dataset consisting of n samples $\mathbf{T}_{1:n}$ of i.i.d. Gaussian tensors of order k with the promise that $\mathbb{E}\mathbf{T}_1$ is a rank-1 tensor and $\|\mathbb{E}\mathbf{T}_1\| = 1$. The goal is to estimate $\mathbb{E}\mathbf{T}_1$. This problem exhibits a large conjectured hard phase when $k > 2$: When $d \lesssim n \ll d^{\frac{k}{2}}$ it is information theoretically possible to estimate $\mathbb{E}\mathbf{T}_1$, but no polynomial time estimator is known. We provide a sharp analysis of the optimal sample complexity in the Statistical Query (SQ) model and show that SQ algorithms with polynomial query complexity not only fail to solve Tensor PCA in the conjectured hard phase, but also have a strictly sub-optimal sample complexity compared to some polynomial time estimators such as the Richard-Montanari spectral estimator. Our analysis reveals that the optimal sample complexity in the SQ model depends on whether $\mathbb{E}\mathbf{T}_1$ is symmetric or not. For symmetric, even order tensors, we also isolate a sample size regime in which it is possible to test if $\mathbb{E}\mathbf{T}_1 = \mathbf{0}$ or $\mathbb{E}\mathbf{T}_1 \neq \mathbf{0}$ with polynomially many queries but not estimate $\mathbb{E}\mathbf{T}_1$. Our proofs rely on the Fourier analytic approach of Feldman et al. (2018) to prove sharp SQ lower bounds.

Keywords: Tensor PCA, Computational-Statistical Gaps, Statistical Query (SQ) Model, Computational Complexity, Matrix/Tensor Estimation.

1. Introduction

In the Tensor PCA testing problem, one is given a dataset consisting of n i.i.d. tensors $\mathbf{T}_{1:n} \in \bigotimes^k \mathbb{R}^d$, $k \geq 2$, and the goal is to determine if the tensors were drawn from the null hypothesis D_0 ,

$$D_0 : (T_i)_{j_1, j_2, \dots, j_k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1) \quad \forall j_1, j_2, \dots, j_k \in [d], \quad \forall i \in [n],$$

or from a distribution in the composite alternate hypothesis,

$$\mathcal{D} = \{D_{\mathbf{V}} : \mathbf{V} \in \bigotimes^k \mathbb{R}^d, \|\mathbf{V}\| = \sqrt{d^k}, \text{Rank}(\mathbf{V}) = 1\},$$

where under the alternate hypothesis $D_{\mathbf{V}}$ the entries of the tensor are distributed as

$$D_{\mathbf{V}} : (T_i)_{j_1, j_2, \dots, j_k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(\frac{V_{j_1, j_2, \dots, j_k}}{d^{\frac{k}{2}}}, 1\right) \quad \forall j_1, j_2, \dots, j_k \in [d], \forall i \in [n].$$

In the Tensor PCA estimation problem, given a dataset $\mathbf{T}_{1:n}$ sampled i.i.d. from an unknown distribution $D \in \mathcal{D}$, the goal is to estimate the *signal tensor* $\mathbb{E}_D \mathbf{T} = \mathbf{V}/\sqrt{d^k}$.

When $k = 2$, this model is called the Spiked Wigner Model and is often attributed to Johnstone et al. (2001), who introduced a variant of this model as a non-null model for Principal Component Analysis. For general k , this model was introduced by Richard and Montanari (2014)¹, who envisioned this model as stylized setup to study statistical and computational phenomena occurring in higher order moment based estimators for fitting latent variable models (Anandkumar et al., 2014). In these applications, a suitable sample moment tensor can also be decomposed as a sum of a low rank population tensor, which contains information about parameters of interest, and a random tensor arising due to finite sample fluctuations.

The most striking feature of the Tensor PCA problem is that it exhibits a large gap between the information theoretically optimal performance and the performance achieved by computationally efficient estimators when $k > 2$. Estimating the mean of a Gaussian tensor in general requires $n \asymp d^k$ samples. However, in the Tensor PCA problems since the mean is low rank, the information theoretic sample complexity requirement is $n \asymp d$. For example, when the signal tensor is symmetric, it is known that there exists a critical value β_k^* such that if $n < (\beta_k^* - o_d(1)) \cdot d$, then it is information theoretically impossible to solve the Tensor PCA testing or estimation problems (Montanari et al., 2015). On the other hand, the results of Richard and Montanari (2014) imply the existence of another threshold β_k^\sharp such that if $n > (\beta_k^\sharp + o_d(1)) \cdot d$, then tests and estimators based on the maximum likelihood value:

$$\max_{\|\mathbf{x}\|=1} \overline{\mathbf{T}}(\mathbf{x}, \mathbf{x}, \dots, \mathbf{x}), \quad \overline{\mathbf{T}} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mathbf{T}_i, \quad (1)$$

have a non-trivial testing and estimation performance. When $k > 2$, no computationally efficient algorithm to solve the optimization problem in Eq. 1 is known. Richard and Montanari have proposed a computationally tractable relaxation of the optimization problem that flattens the mean tensor $\overline{\mathbf{T}}$ into a $d^{\lceil \frac{k}{2} \rceil} \times d^{\lfloor \frac{k}{2} \rfloor}$ matrix and uses the largest singular value and vector of the resulting matrix for testing and estimation respectively. This procedure requires $n \gtrsim d^{\frac{k}{2}}$ to solve the testing and estimation problems (Hopkins et al., 2015; Zheng and Tomioka, 2015). To this date no known computationally efficient procedure improves on the performance of the Richard-Montanari spectral method. Consequently it is conjectured that any computationally efficient testing or estimation procedure requires $n \gtrsim d^{\frac{k}{2}}$ samples. Proving this conjecture is expected to be more difficult than resolving P v.s. NP since proving computational lower bounds for average-case problems is harder than proving

1. In the standard formulation of Tensor PCA introduced by Richard and Montanari (2014) one only observes the empirical average of the tensors $\overline{\mathbf{T}}$. However this is equivalent to the model studied here since given the n samples $\mathbf{T}_{1:n}$, one can compute $\overline{\mathbf{T}}$, and given $\overline{\mathbf{T}}$, one can generate an i.i.d. sample $\mathbf{T}_{1:n}$ without knowing $\mathbb{E}\mathbf{T}$ since $\overline{\mathbf{T}}$ is a sufficient statistic for $\mathbb{E}\mathbf{T}$.

computational lower bounds for worst-case problems. However, recent works have exhibited some evidence that lends credibility to this conjecture.

In this paper, we study lower bounds on a class of algorithms for Tensor PCA: those that can be implemented in Statistical Query (SQ) computational model introduced by Kearns (1998). In the SQ Model, the estimation or testing procedure does not have direct, unrestricted access to the data $\mathbf{T}_{1:n}$ generated from the unknown distribution $D \in \mathcal{D} \cup \{D_0\}$, but instead can access the data only through an oracle. The algorithm can query the oracle with a statistic q of interest, and the oracle returns an estimate of the expected value of the statistic $\mathbb{E}_D q(\mathbf{T})$. The response of the oracle, denoted by $\hat{\mathbb{E}}_D q(\mathbf{T})$, has an error guarantee comparable to the error guarantee of the mean of the statistic on the dataset, but can be otherwise arbitrary (see Definition 1); that is,

$$\left| \hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T}) \right| \approx \left| \frac{1}{n} \sum_{i=1}^n q(\mathbf{T}_i) - \mathbb{E}_D q(\mathbf{T}) \right|. \quad (2)$$

The SQ framework to study computational-statistical gaps aims to show that in the regime where no polynomial time estimators are known for the learning problem, no SQ algorithm can solve the learning problem after making only polynomially many queries. We emphasize that a successful SQ algorithm must work with an arbitrary oracle whose errors $|\hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T})|$ are possibly adversarial but do satisfy the requirement of Eq. 2, and not just the most natural sample mean oracle which estimates $\mathbb{E}_D q(\mathbf{T})$ by the empirical mean of $q(\cdot)$ on n samples drawn from D . This strong robustness requirement which is expected in the SQ model is what makes it possible to prove strong lower bounds in this model. At the same time, one hopes that though most practical procedures require direct access to the samples $\mathbf{T}_{1:n}$, they are essentially computing certain adaptively chosen statistics q on the dataset. Moreover, one hopes that their analysis only relies on the fact that the sample mean of these statistics satisfies an error guarantee of the form given in Eq. 2. It would then be possible to implement these procedures in the SQ model without any degradation in their performance and hence, lower bounds proved in the SQ model would apply to these procedures.

Indeed, these hopes have been realized in a number of problems which exhibit statistical-computational gaps. Two notable examples are the Planted Clique problem (Feldman et al., 2017a) and sparse PCA (Wang et al., 2015). Both of these problems exhibited a gap between the information theoretically optimal performance and the performance of the best known polynomial time estimator. The performance of the best known polynomial time estimator can be achieved in the SQ model and moreover is the best possible performance in the SQ model. Hence, for these problems the SQ Model provides a framework to prove lower bounds on the sample complexity of a class of computationally tractable estimators which include the best known estimators for this problem. Furthermore, the SQ lower bounds for these problems provide an intuitive *mechanistic explanation* for the observed statistical-computational gap based on the informational structure of the problem: In order to declare that the unknown distribution $D \notin \mathcal{D}$, an SQ algorithm must rule out all alternate hypothesis $D_1 \in \mathcal{D}$ by asking queries such that $|\hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_{D_1} q(\mathbf{T})|$ is more than the finite sample fluctuations expected from n samples. The SQ lower bound proceeds by exhibiting an *exponentially* large (in d) collection of distributions $\mathcal{H} \subset \mathcal{D}$ such that even the single

most powerful query is able to rule out only a small fraction of the distributions in \mathcal{H} (more precisely, $o(d^{-t})$ fraction for any $t \in \mathbb{N}$). Consequently, any SQ algorithm must make a super-polynomial number of queries to solve these problems in the regimes where they are believed to be computationally hard. More recently, SQ lower bounds have also been used to reason about computational hardness of various robust estimation problems (Diakonikolas et al., 2017, 2019). We also refer the reader to the recent review article on the SQ model (Reyzin, 2020).

1.1 Problem Definition and the Statistical Query Model

Recall that in the alternative hypothesis for Tensor PCA, $\mathbb{E}\mathbf{T} = \mathbf{V}/\sqrt{d^k}$, where \mathbf{V} is a rank-1 tensor with $\|\mathbf{V}\| = \sqrt{d^k}$. Hence we can write \mathbf{V} as:

$$\mathbf{V} = \mathbf{v}_1 \otimes \mathbf{v}_2 \cdots \otimes \mathbf{v}_k, \mathbf{v}_i \in \mathbb{R}^d, \|\mathbf{v}_i\| = \sqrt{d} \forall i \in [k].$$

In the standard symmetric version of Tensor PCA introduced by Richard and Montanari, one further assumes $\mathbf{v}_1 = \mathbf{v}_2 = \cdots = \mathbf{v}_k$. However, the completely asymmetric version where this constraint is not enforced is also common in the literature (Zheng and Tomioka, 2015; Zhang and Xia, 2018). It turns out that the asymmetric Tensor PCA problems are harder to solve for SQ algorithms that make only polynomially many queries. In order to get more insight into this phenomena, in this paper, we study a generalization of Tensor PCA to incorporate partial symmetries where only some of the \mathbf{v}_i 's are identical. This interpolates between symmetric Tensor PCA on one end and completely asymmetric Tensor PCA on the other end. The generalization is parametrized by integers $k, K \in \mathbb{N}$ with $K \leq k$ and a labelling function $\pi : [k] \rightarrow [K]$ such that, under the alternate hypothesis,

$$\mathbb{E}\mathbf{T} = \frac{1}{\sqrt{d^k}} \cdot \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \otimes \cdots \otimes \mathbf{v}_{\pi(k)}.$$

For example, if $k = 3, K = 2$ and $\pi(1) = \pi(2) = 1, \pi(3) = 2$, then $\sqrt{d^3} \cdot \mathbb{E}\mathbf{T} = \mathbf{v}_1 \otimes \mathbf{v}_1 \otimes \mathbf{v}_2$, for two, possibly distinct $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^d$ with $\|\mathbf{v}_1\| = \|\mathbf{v}_2\| = \sqrt{d}$. Hence, we define the composite alternative hypothesis:

$$\mathcal{D}(\pi) = \{\mathcal{D}_{\mathbf{V}} : \mathbf{V} = \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \otimes \cdots \otimes \mathbf{v}_{\pi(k)}, \mathbf{v}_{1:K} \in \mathbb{R}^d, \|\mathbf{v}_i\| = \sqrt{d} \forall i \in [K]\}.$$

In order to prove the strongest possible lower bounds, we will assume the algorithm knows the labelling function π . In the (generalized) Tensor PCA testing problem, the testing algorithm is given data $\mathbf{T}_{1:n}$ sampled i.i.d. from some unknown $D \in \{D_0\} \cup \mathcal{D}(\pi)$ and has to decide if $D = D_0$ or $D \in \mathcal{D}(\pi)$. In the (generalized) Tensor PCA estimation problem, the data is generated from an unknown $D \in \mathcal{D}(\pi)$ and one seeks to estimate $\mathbb{E}_D \mathbf{T}$ with small error.

We study tests and estimators for Tensor PCA in the Statistical Query framework, where the algorithm can access the data only through the VSTAT oracle introduced by Feldman et al. (2017a).

Definition 1 (VSTAT oracle for Tensor PCA) For any $D \in \{D_0\} \cup \mathcal{D}(\pi)$ and any $n \in \mathbb{N}$, $\text{VSTAT}_D(n)$ is an oracle that, when given as input an arbitrary function $q :$

$\otimes^k \mathbb{R}^d \mapsto [0, 1]$, returns an estimate $\hat{\mathbb{E}}_D q(\mathbf{T})$ of the expectation $\mathbb{E}_D q(\mathbf{T})$ with the error guarantee:

$$|\hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T})| \leq \max \left(\frac{1}{n}, \sqrt{\frac{\mathbb{E}_D q(\mathbf{T}) \cdot (1 - \mathbb{E}_D q(\mathbf{T}))}{n}} \right). \quad (3)$$

The parameter n is known as the effective sample size of the oracle.

The rationale for defining the oracle in this way is that the error guarantee of the oracle given in Eq. 3 is comparable to the error guarantee obtained if one estimated $\mathbb{E}_D q(\mathbf{T})$ with its sample average on a dataset of size n drawn from the distribution D . This is also why the parameter n in $\text{VSTAT}_D(n)$ is interpreted as the effective sample size. Our goal is to understand how large must the effective sample size n be for there to exist an SQ algorithm to solve the Tensor PCA testing problem by making only polynomial number of queries in the following sense:

Definition 2 (Tensor PCA Testing in the SQ Model) *An algorithm \mathcal{A} solves the Tensor PCA testing problem in the SQ model with n samples and B queries if, for any $D \in \{D_0\} \cup \mathcal{D}(\pi)$ (unknown to \mathcal{A}), given access to an arbitrary $\text{VSTAT}_D(n)$ oracle, \mathcal{A} correctly decides if $D = D_0$ or $D \in \mathcal{D}(\pi)$ after making at most B (possibly adaptive) queries to $\text{VSTAT}_D(n)$.*

Definition 3 (Tensor PCA Estimation in the SQ Model) *An algorithm \mathcal{A} solves the Tensor PCA estimation problem in the SQ model with n samples and B queries if, for any $D \in \mathcal{D}(\pi)$ (unknown to \mathcal{A}), given access to an arbitrary $\text{VSTAT}_D(n)$ oracle, \mathcal{A} returns an estimate $\hat{\mathbf{V}}$ such that*

$$\|\hat{\mathbf{V}} - \mathbb{E}_D \mathbf{T}\| \leq \frac{1}{4},$$

after making at most B (possibly adaptive) queries to $\text{VSTAT}_D(n)$.

The sample complexity of Tensor PCA in the Statistical Query depends on some structural properties of the labelling function π . For any $i \in [K]$, let s_i be the number of times label i is used in the labelling function π :

$$s_i \stackrel{\text{def}}{=} |\pi^{-1}(i)| = |\{j \in [k] : \pi(j) = i\}|. \quad (4)$$

The sample complexity of Tensor PCA depends crucially on the number of labels that are used an odd number of times. We denote this parameter by \circ and call it the *oddness parameter* of the labelling function π , i.e., $\circ = |\{i \in [K] : s_i \text{ is odd}\}|$. Note that $\circ \in \{0, 1, \dots, k\}$.

1.2 Our Contribution

Our main result is the following SQ lower bound for Tensor PCA.

Theorem 4 *Let $C_0 \geq 0$ and $\epsilon \in (0, 1)$ be arbitrary constants. Consider the learning problems:*

1. Tensor PCA testing with $\circ = 0$ and sample size $n \leq C_0 d^{\frac{k}{2}-\epsilon}$.
2. Tensor PCA estimation with $\circ = 0$ and sample size $n \leq C_0 d^{\frac{k+2}{2}-\epsilon}$.
3. Tensor PCA testing or estimation with $\circ \geq 1$ and sample size $n \leq C_0 d^{\frac{k+\circ}{2}-\epsilon}$.

Then for any $L \in \mathbb{N}$, there exists a positive constant $c_1(k, L, \epsilon, C_0) > 0$ depending only on k, L, ϵ, C_0 and a finite constant $C_1(k, L, \epsilon) < \infty$ depending only on k, L, ϵ such that for all $d \geq C_1(k, L, \epsilon)$, any SQ algorithm that solves any of the above problems must make at least $c_1(k, L, \epsilon, C_0) \cdot d^L$ queries.

We also design optimal SQ procedures, which show that the above lower bounds are tight.

Theorem 5 *There exist SQ algorithms that make $O(d^k \cdot \log(n))$ queries and solve:*

1. The Tensor PCA testing problem with $\circ = 0$ provided $n \gg \sqrt{d^k}$.
2. The Tensor PCA estimation problem with $\circ = 0$ provided $n \gg \sqrt{d^{k+2}}$.
3. The Tensor PCA testing and estimation problems with $\circ \geq 1$ provided $n \gg \sqrt{d^{k+\circ}}$.

Our results have the following implications:

1. As mentioned previously, when $k \geq 3$, no known computationally efficient algorithm is able to leverage the low rank structure in Tensor PCA optimally. Our results show that SQ algorithms also do not work in the conjectured hard phase $d \lesssim n \ll \sqrt{d^k}$.

The SQ framework leverages the adversarial corruptions in the query responses to prove strong lower bounds. However, our result shows that these adversarial corruptions have the unintended effect of making the SQ lower bounds too pessimistic:

2. When $k = 2$, the estimator based on the best rank-1 approximation to $\bar{\mathbf{T}}$ is able to effectively leverage the low dimensional structure and attain the information theoretic sample complexity in polynomial time. In contrast, SQ estimators with polynomial query complexity are ineffective at taking advantage of this low dimensional structure and require as many samples as in the case when the mean is unstructured ($n \gtrsim d^2$).
3. When $k \geq 3$, SQ estimators with polynomially many queries are still less effective than known polynomial time algorithms in taking advantage of the low rank structure, however now, the effectiveness of SQ procedures depends crucially on the underlying symmetry of the mean tensor. They are least effective when the mean tensor is completely asymmetric, where they require as many samples as in the unstructured case ($n \gtrsim d^k$).
4. When $\circ = 0$ (e.g., even-order symmetric Tensor PCA), the problem exhibits an estimation-testing gap in the SQ model: in the regime $\sqrt{d^k} \lesssim n \ll \sqrt{d^{k+2}}$ testing is possible but estimation is not. No such testing-estimation gap is seen outside the SQ model.

The ineffectiveness of SQ algorithms seems connected to the fact that leveraging the low rank structure requires solving stochastic non-convex optimization with adversarial, inexact gradient oracles. The results of Feldman et al. (2017b) show that stochastic *convex* optimization problems based on n samples (such as the least squares problem) can be solved with restricted access to a $\text{VSTAT}(n)$ oracle with no degradation in performance. In contrast, our estimation lower bound for the symmetric spiked Wigner problem ($k = 2, \mathbf{o} = 0, \mathbb{E}\mathbf{T} = \mathbf{v}_0\mathbf{v}_0^\top$) exhibits a canonical, benign stochastic non-convex optimization problem: $\max_{\|\mathbf{v}\|=1} \mathbb{E}\langle \mathbf{v}, \mathbf{T}\mathbf{v} \rangle$ which cannot be solved with a $\text{VSTAT}(n)$ oracle without a degradation in sample complexity compared to its empirical version: $\max_{\|\mathbf{v}\|=1} \langle \mathbf{v}, \overline{\mathbf{T}}\mathbf{v} \rangle$. This rules out non-convex extensions to the results of Feldman et al. (2017b) without a worse dependence on problem parameters. This suboptimality of SQ algorithms seems to arise because random initialization starts close to a saddle point of the population objective. In the spiked Wigner case, the n sample empirical gradient at initialization \mathbf{v}_{init} is $2\langle \mathbf{v}_0, \mathbf{v}_{\text{init}} \rangle \mathbf{v}_0 + 2\mathbf{g}$ where $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d/n)$. Although the “signal” part is small, $\langle \mathbf{v}_0, \mathbf{v}_{\text{init}} \rangle \approx d^{-\frac{1}{2}}$, the noise \mathbf{g} is randomly oriented, $\langle \mathbf{g}, \mathbf{v}_0 \rangle \approx n^{-1/2}$, and hence when $n \gg d$, the signal part dominates the noise part in the projection of the gradient along \mathbf{v}_0 at initialization. This signal is amplified in the subsequent iterations of the power method. In contrast, when a $\text{VSTAT}(n)$ oracle is used, the approximate gradient is of the form $2\langle \mathbf{v}_0, \mathbf{v}_{\text{init}} \rangle \mathbf{v}_0 + 2\tilde{\mathbf{g}}$, where $\|\tilde{\mathbf{g}}\|^2 \approx \|\mathbf{g}\|^2 \approx d/n$; but the noise vector $\tilde{\mathbf{g}}$ is no longer randomly oriented and can be adversarially chosen to cancel out the signal unless $n \gg d^2$.

Given that SQ algorithms are suboptimal for Tensor PCA, one might wonder why one should care about SQ lower bounds for Tensor PCA. Indeed, the SQ prediction of the threshold for hardness of Tensor PCA is nowhere near the correct (conjectured) threshold, so SQ lower bounds are not evidence for the computational hardness of the problem. We believe that SQ lower bounds for Tensor PCA are still useful for two reasons. (1) Currently, there are several frameworks to predict computational-statistical gaps including SQ lower bounds, Hopkins’ Low Degree method (Hopkins, 2018; Kunisky et al., 2019), reductions from other hard problems (Berthet and Rigollet, 2013), and statistical physics heuristics like failure of Belief Propagation (Zdeborová and Krzakala, 2016). We believe it is of interest to report successes and failures of these frameworks in canonical problems to understand their relative strengths and weaknesses. (2) SQ lower bounds are also useful even beyond the study of computational-statistical gaps. SQ algorithms are robust since they can’t use specialized properties of finite sample fluctuations implied by model assumptions. This is desirable in practice, and it is of interest to understand how the sample complexity changes under some natural deviations from model assumptions.

1.3 Related Work

A number of works have investigated the computational-statistical gap in Tensor PCA, which we discuss below.

Studies of Computational-Statistical Gap in Tensor PCA: Ben Arous et al. (2019) have shown that in the information theoretic regime $n \asymp d$, the expected number of local maximizers of the maximum likelihood optimization problem (Eq. 1) is exponentially large. These local maximizers can potentially trap gradient based methods. Ben Arous et al. (2018) have shown that when $n \gg d^{k-1}$, Langevin Dynamics succeeds in solving Tensor

PCA and that this is essentially tight for Langevin Dynamics to work. Sum of Squares relaxations for Tensor PCA have been analyzed (Hopkins et al., 2015, 2016, 2017; Bhatiprolu et al., 2016) and it is known that computationally efficient sum-of-squares based procedures match the performance of the Richard-Montanari spectral estimator but fail to improve on it. A popular way to provide evidence for hardness of a particular problem is via reductions, i.e., to show that solving it would lead to algorithms to solve another problem which is widely believed to be hard. Reductions were introduced to reason about the hardness of statistical inference problems in the seminal work of Berthet and Rigollet (2013). In the specific case of Tensor PCA, it has been shown that hardness of Hypergraph Planted Clique implies the hardness of Tensor PCA (Zhang and Xia, 2018; Brennan and Bresler, 2020).

Our results reveal a gap between the optimal sample complexity of Tensor PCA for polynomial query SQ algorithms and computationally efficient algorithms. A well known example of this situation is efficient learning of parity functions (in a noise-free setting), which is possible using the Gaussian elimination algorithm, but not using SQ algorithms (Blum et al., 1994). Apart from this, we are aware of at least two other instances where SQ algorithms are suboptimal in comparison to some computationally efficient and noise tolerant procedure, which we discuss next.

Planted Satisfiability: Feldman et al. (2018) study a broad class of planted random constraint satisfaction problems in which an unknown assignment $\sigma_\star \in \{0, 1\}^d$ is fixed and we observe n clauses C_1, C_2, \dots, C_n drawn independently and uniformly at random from a family of clauses \mathcal{C} , conditioned on σ_\star being a satisfying assignment for the clause. The goal is to identify the planted assignment σ_\star . When \mathcal{C} is the set of all XOR expressions involving k variables, the problem is called the planted k -XOR-SAT problem and exhibits a large computational-statistical gap: observing $n \gtrsim d \log(d)$ clauses is information theoretically sufficient to identify the planted assignment σ_\star with high probability. However, all noise-tolerant and computationally efficient methods require $n \gtrsim d^{\frac{k}{2}}$ clauses. On the other hand, Feldman et al. (2018) show that SQ algorithms fail unless $n \gtrsim d^k$. Moreover, this work introduces the Fourier analytic approach to prove SQ lower bounds that we rely on. Previous approaches for proving SQ lower bounds (Feldman et al., 2017a) are only able to show a weaker sample complexity lower bounds of $n \gtrsim d^{\frac{k}{2}}$ for k -XOR-SAT as well as k -Tensor PCA. This work also proposes a hierarchical generalization of the VSTAT(n) oracle for satisfiability problems known as the MVSTAT(n, ℓ) oracle where ℓ is a parameter controlling the strength of the oracle. The oracle is designed such that, for a well-chosen value of ℓ , it is possible to implement a particular $d^{\frac{k}{2}}$ sample complexity estimator for k -XOR-SAT without any degradation in performance.

Matrix Mean Estimation: The recent work Li et al. (2019) studies mean estimation on general normed spaces. Their results on matrix mean estimation with respect to the operator norm ($\|\cdot\|_{\text{op}}$) are particularly relevant for our problem. In this problem, one is given n samples $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ of $d \times d$ matrices sampled i.i.d. from some distribution supported on unit operator norm ball, and the goal is to estimate $\mathbb{E}\mathbf{X}$ with an error of $\epsilon = 0.1$ in operator norm. By the matrix Bernstein inequality (Tropp et al., 2015), when $n \gtrsim \log(d)$, the empirical average of the matrices $\bar{\mathbf{X}}$ satisfies $\|\bar{\mathbf{X}} - \mathbb{E}\mathbf{X}\|_{\text{op}} \leq 0.1$. On the other hand, Li et al. show that in the SQ model one requires $n \gtrsim d$ samples to solve this problem. Though their results are stated for unstructured matrix mean estimation, the

hard distribution they construct is supported on highly sparse Boolean matrices with a low rank expectation. Thus, after a suitable change of scale, their hard instance also shows $n \gtrsim d^2$ sample complexity lower bound in the SQ model for a non-Gaussian variant of low rank matrix mean estimation. Their proof also relies on the Fourier analytic approach of Feldman et al. (2018).

Another commonly used framework to provide evidence for a computational-statistical gap is the low-degree polynomial framework, which we discuss next.

Connection with low-degree polynomial lower bounds: The low-degree polynomial framework seeks to provide evidence for a computational-statistical gap by showing that any procedure that seeks to solve the inference problem by computing a low-degree polynomial of the data fails to solve the inference in the conjectured hard phase (Hopkins, 2018; Kunisky et al., 2019). This framework consistently reproduces widely believed average-case hardness conjectures for many problems, including Tensor PCA (Hopkins et al., 2017; Hopkins, 2018; Kunisky et al., 2019). The concurrent work of Brennan et al. (2020) identifies general conditions under which lower bounds for low-degree polynomial procedures are equivalent to a *sufficient criterion* for SQ lower bounds, namely the statistical dimension criterion of Feldman et al. (2017a). Hence, for many inference problems, including the completely symmetric Tensor PCA testing problem, they obtain SQ lower bounds from previously known low-degree lower bounds. However, the SQ lower bounds obtained in this manner are often not tight. For example, Brennan et al. show that when $n \ll d^{\frac{3}{2}}$, SQ algorithms with polynomial query complexity fail to solve the completely symmetric Tensor PCA testing problem with $k = 3$. In contrast, our results show that the same conclusion holds when $n \ll d^2$. The reason for this appears to be that the statistical dimension criterion of Feldman et al. (2017a) studied by Brennan et al. is only a sufficient (and not necessary) criterion for SQ lower bounds. In contrast, our lower bounds rely on a different criterion introduced by Feldman et al. (2018) which is known to imply stronger SQ lower bounds in other problems like planted CSPs (Feldman et al., 2018).

1.4 Notation

In this section, we introduce the notation used in this paper.

Common Sets: $\mathbb{R}, \mathbb{N}, \mathbb{N}_0$ denote the set of real numbers, natural numbers and non-negative integers respectively. We denote the finite set $\{1, 2, \dots, k\}$ by $[k]$. The set $\otimes^k \mathbb{R}^d$ denotes the set of all $d \times d \times \dots \times d$ tensors with real entries. The set $\otimes^k \mathbb{N}_0^d$ denotes the set of all $d \times d \times \dots \times d$ tensors with non-negative integer entries.

Vectors, Matrices and Tensors: We denote vectors, matrices and tensors with bold face letters. We denote the d dimension vectors $(1, 1, \dots, 1)$ and $(0, 0, \dots, 0)$ as $\mathbf{1}_d$ and $\mathbf{0}_d$. We will omit the subscript d when it is clear from the context. For vectors, $\|\cdot\|$ and $\|\cdot\|_1$ denote the Euclidean (ℓ_2) norm and the ℓ_1 norm respectively. For vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$, $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the standard Euclidean inner product $\sum_{i=1}^d a_i b_i$. These norms and inner product are analogously defined for matrices and tensors by stacking their entries into a single vector. For a tensor $\mathbf{T} \in \otimes^k \mathbb{N}_0^d$, we define the map $\text{SUM}(\mathbf{T}, l) : \otimes^k \mathbb{N}_0^d \times [k] \mapsto \mathbb{R}^d$ which computes the partial sum of the entries of \mathbf{T} along all modes except mode l :

$$\text{SUM}(\mathbf{T}, l)_i = \sum_{j_1, j_2, \dots, j_{l-1}, j_{l+1}, \dots, j_k \in [d]} T_{j_1, j_2, \dots, j_{l-1}, i, j_{l+1}, \dots, j_k}. \quad (5)$$

Likewise we define $\text{PARITY}(\cdot, \cdot) : \otimes^k \mathbb{N}_0^d \times [k] \mapsto \{0, 1\}^d$ as follows:

$$\text{PARITY}(\mathbf{T}, l)_i = \bigoplus_{j_1, j_2, \dots, j_{l-1}, j_{l+1}, \dots, j_k \in [d]} T_{j_1, j_2, \dots, j_{l-1}, i, j_{l+1}, \dots, j_k}. \quad (6)$$

In the above display, \oplus denotes addition modulo 2.

Given a labelling function $\pi : [k] \mapsto [K]$, which assigns each mode $l \in [k]$ to a label $\pi(l) \in K$, we can combine the partial sums which leave out modes with the same label by adding them. This defines a map $\text{SUM}_\pi(\mathbf{T}, i) : \otimes^k \mathbb{N}_0^d \times [K] \mapsto \mathbb{R}^d$:

$$\text{SUM}_\pi(\mathbf{T}, i) = \sum_{l: \pi(l)=i} \text{SUM}(\mathbf{T}, l).$$

Analogously, we combine the partial parities which leave out modes with the same label via the map $\text{PARITY}_\pi(\mathbf{T}, i) : \otimes^k \mathbb{N}_0^d \times [K] \mapsto \{0, 1\}^d$ defined as follows:

$$\text{PARITY}_\pi(\mathbf{T}, i) = \bigoplus_{l: \pi(l)=i} \text{PARITY}(\mathbf{T}, l).$$

For a tensor \mathbf{T} , the support of \mathbf{T} is $\text{Supp}(\mathbf{T}) = \{(i_1, i_2, \dots, i_k) \in \mathbb{N}^k : T_{i_1, i_2, \dots, i_k} \neq 0\}$.

Special Distributions: $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. $\text{Pois}(\lambda)$ denotes the Poisson distribution with rate $\lambda \geq 0$. $\text{Mult}(n, k)$ denotes the Multinomial distribution governing the random vector that counts the occurrences of $1, 2, \dots, k$ when a fair k -faced die is independently thrown n times. For a finite set S , $\text{Unif}(S)$ denotes the uniform distribution on S , and $\text{Unif}(\{\pm 1\})$ is also called the Rademacher distribution.

Asymptotic Order Notation: We use the following notation to describe the behaviour of positive sequence $f(d)$ and $g(d)$ for large d . We say $f(d) \lesssim g(d)$ if there exists a constant c independent of d such that $f(d) \leq cg(d)$. If $f(d) \lesssim g(d)$ and $g(d) \lesssim f(d)$, then we say $f(d) \asymp g(d)$. Finally, we say a sequence $f(d) \ll g(d)$ if there exists an $\epsilon > 0$ such that $f(d)d^\epsilon \lesssim g(d)$.

Miscellaneous: For real numbers a, b , $a \vee b$ denotes the maximum of a, b . For a vector $\mathbf{v} \in \mathbb{R}^d$ and a vector $\mathbf{c} \in \mathbb{N}_0^d$, we define the entry-wise factorial and the entry-wise powering notations:

$$\mathbf{c}! \stackrel{\text{def}}{=} \prod_{i=1}^d c_i!, \quad \mathbf{a}^{\mathbf{c}} \stackrel{\text{def}}{=} \prod_{i=1}^d a_i^{c_i}. \quad (7)$$

We also use the natural extension of the above entry-wise operators to the case when \mathbf{a}, \mathbf{c} are tensors.

2. Proof Overview for Theorem 4

For Tensor PCA with the labelling function π , we consider the following prior on $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$: $\mathbf{v}_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\}^d)$. Define the random tensor $\mathbf{V} = \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \otimes \dots \otimes \mathbf{v}_{\pi(k)}$. Note that for any realization of \mathbf{V} , $D_{\mathbf{V}} \in \mathcal{D}(\pi)$. Standard results on the SQ framework (developed in Feldman et al. (2018), reviewed in Appendix A) show that the minimal possible query

complexity of an SQ algorithm that solves the Tensor PCA testing problem with n samples is determined by the asymptotics of the following quantity:

$$\max_{q: \mathbb{E}_{D_0} q^2(\mathbf{T})=1} \Delta(q; D_0), \text{ where } \Delta(q; D_0) \stackrel{\text{def}}{=} \mathbb{P}_{\mathbf{V}} \left[|\mathbb{E}_{D_{\mathbf{V}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| \gtrsim \frac{1}{\sqrt{n}} \right]. \quad (8)$$

The intuition for this is as follows. Consider a fixed realization of the r.v. \mathbf{V} . Suppose for a query q , we have $|\mathbb{E}_{D_{\mathbf{V}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| \gtrsim n^{-\frac{1}{2}}$. Then given a response from an arbitrary $\text{VSTAT}_D(n)$ oracle, one can distinguish if $D = D_0$ or if $D = D_{\mathbf{V}}$ since the gap $|\mathbb{E}_{D_{\mathbf{V}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})|$ is more than the size of the error introduced by the $\text{VSTAT}_D(n)$ oracle. In this situation, we say that the query q has *resolved* the alternative hypothesis $D_{\mathbf{V}}$. In order to solve the composite hypothesis testing problem of distinguishing $D = D_0$ or $D \in \mathcal{D}(\pi)$, one must ask queries to resolve at least every hypothesis $D_{\mathbf{V}}$ for every possible realization of \mathbf{V} . Note that for a query q , the quantity $\Delta(q; D_0)$ measures the information content of a single query: the fraction of realizations of \mathbf{V} that are resolved by the query q . In order to minimize the number of queries asked, one should ask powerful queries: single queries which resolve a large fraction of hypotheses in one go. Hence, in order to show that no SQ algorithm can solve Tensor PCA with polynomially many queries it is sufficient to show that any single query resolves at most $o(d^{-t})$ fraction of possible realizations of $\mathbf{V} \forall t \in \mathbb{N}$.

For the estimation problem, the results of Feldman (2017b, Theorem 4.2) show that the number of queries that need to be made to solve the estimation problem with n samples is governed by the asymptotics of a quantity analogous to object in Eq. 8 with a small change: the measure D_0 can be replaced by an arbitrary measure on D_{\star} on $\bigotimes^k \mathbb{R}^d$. For most cases, choosing $D_{\star} = D_0$ suffices. In other cases we will choose

$$D_{\star} = \overline{D}(\pi) \stackrel{\text{def}}{=} D_{\overline{\mathbf{V}}(\pi)}, \text{ where } \overline{\mathbf{V}}(\pi) \stackrel{\text{def}}{=} \mathbb{E} \mathbf{V}. \quad (9)$$

To analyze Eq. 8, we use a key idea of Feldman et al. (2018): $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \otimes \dots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$ is a function of Boolean random vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$, and Hypercontractivity Theorems can be used to control its moments and concentration properties. Implementing this approach requires an understanding of the Fourier spectrum of $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \otimes \dots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$, which is possible due to the additive Gaussian structure of Tensor PCA. This structure is also leveraged in the Low Degree method for predicting computational-statistical gaps (Hopkins, 2018; Kunisky et al., 2019). This results in the following proposition, which is proved in Appendix B.

Proposition 6 *Let $\epsilon > 0$, $u \in \mathbb{N}, u \geq 2$ be arbitrary. For the reference distribution D_0 , and for any query q with $\mathbb{E}_{D_0} q^2(\mathbf{T}) = 1$, we have,*

$$\mathbb{P}_{\mathbf{V}} [|\mathbb{E}_{D_{\mathbf{V}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \epsilon] \leq \left(\frac{\epsilon}{\epsilon^2} \cdot \max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left((u-1)^{l_1+l_2+\dots+l_K} \cdot p_{\pi}(l_{1:K}) \right) \right)^{\frac{u}{2}},$$

where the coefficients $p_{\pi}(l_1, l_2, \dots, l_K)$ are defined as follows:

$$p_{\pi}(l_1, l_2, \dots, l_K) \stackrel{\text{def}}{=} \frac{1}{e} \sum_{\substack{\mathbf{c} \in \bigotimes^k \mathbb{N}_0^d: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_{\pi}(\mathbf{c}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]}} \frac{1}{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}.$$

For the reference measure $\bar{D}(\pi)$, for any query q such that $\mathbb{E}_{\bar{D}(\pi)} q^2(\mathbf{T}) = 1$, we have,

$$\mathbb{P}_{\mathbf{V}} \left[\left| \mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{\bar{D}(\pi)} q(\mathbf{T}) \right| > \epsilon \right] \leq \left(\frac{e}{\epsilon^2} \cdot \max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left((u-1)^{l_1 + l_2 + \dots + l_K} \cdot \bar{p}_\pi(l_{1:K}) \right) \right)^{\frac{u}{2}},$$

where the coefficients $\bar{p}_\pi(l_1, l_2, \dots, l_K)$ are defined as follows:

$$\bar{p}_\pi(l_1, l_2, \dots, l_K) \stackrel{\text{def}}{=} \frac{1}{e} \sum_{\substack{\mathbf{c} \in \otimes^k \mathbb{N}_0^d : \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]}} \frac{1}{d^k \|\mathbf{c}\|_1 \cdot \mathbf{c}!} \cdot \mathbb{I}(\text{Supp}(\mathbf{c}) \cap \text{Supp}(\bar{\mathbf{V}}(\pi)) = \phi).$$

2.1 Asymptotics of Combinatorial Coefficients p_π and \bar{p}_π

Proving sharp SQ lower bounds for Tensor PCA requires a tight asymptotic analysis of the combinatorial coefficients p_π and \bar{p}_π . In order to illustrate why this requires a delicate analysis, we consider the following natural bound:

$$p_\pi(l_1, l_2, \dots, l_K) = \frac{1}{e} \sum_{\substack{\mathbf{c} \in \otimes^k \mathbb{N}_0^d : \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]}} \frac{1}{d^k \|\mathbf{c}\|_1 \cdot \mathbf{c}!} \leq \frac{1}{ed^k} \sum_{\mathbf{c} \in \otimes^k \mathbb{N}_0^d} \frac{1}{\mathbf{c}!} = \frac{1}{ed^k} e^{d^k}, \quad (10)$$

where the last step follows from $e = 1 + 1/1! + 1/2! + \dots$. This estimate diverges exponentially with d^k and is too weak for our analysis. In order to obtain tight estimates, we rely on a natural probabilistic interpretation of the coefficients p_π as a certain Poisson probability. Let \mathbf{C} be a random tensor in $\otimes^k \mathbb{N}_0^d$ whose entries are sampled independently as $C_{i_1, i_2, \dots, i_k} \stackrel{\text{i.i.d.}}{\sim} \text{Pois}(d^{-k})$. Recalling the formula for the probability mass function of the Poisson distribution (Fact 21), we have

$$\begin{aligned} p_\pi(l_1, l_2, \dots, l_K) &\stackrel{\text{def}}{=} \sum_{\substack{\mathbf{c} : \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]}} \prod_{i_1, i_2, \dots, i_k \in [d]} \frac{e^{-\frac{1}{d^k}}}{(d^k)^{c_{i_1, \dots, i_k}} \cdot c_{i_1, \dots, i_k}!} \\ &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]). \end{aligned}$$

Analogously,

$$\begin{aligned} \bar{p}_\pi(l_1, l_2, \dots, l_K) \\ &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\bar{\mathbf{V}}(\pi)) = \phi). \end{aligned}$$

This interpretation already shows that $p_\pi \leq 1$ and $\bar{p}_\pi \leq 1$, which is much better than the trivial estimate in Eq. 10. We obtain tighter bounds by taking into account the parity constraints using the following conditional independence structure in Poisson random tensors.

Proposition 7 *Let \mathbf{C} be a Poisson random tensor sampled as above. Then, we have*

$$\|\mathbf{C}\|_1 \sim \text{Pois}(1).$$

Moreover, conditioned on $\|\mathbf{C}\|_1$, we have

$$\text{SUM}(\mathbf{C}, i) \mid \|\mathbf{C}\|_1 \stackrel{i.i.d.}{\sim} \text{Mult}(\|\mathbf{C}\|_1, d), \forall i \in [k].$$

The proof of this property can be found in Appendix F. The following result provides an exact formula for the probabilities $\mathbb{P}(\text{PARITY}(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K])$ which in principle allows us to compute the complete asymptotic expansion for these probabilities for large d .

Lemma 8 *Let $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_K$ be K independent vectors in $\{\pm 1\}^d$ whose entries $(X_i)_j$ are i.i.d. Rademacher distributed and independent of \mathbf{C} . Let \bar{X}_i denote the mean of the entries of the random vector \mathbf{X}_i . Then,*

$$\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]) = \frac{1}{e} \cdot \mathbb{E} \left[e^{\bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \dots \bar{X}_K^{s_K}} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right], \quad (11)$$

where $s_{1:K}$ were defined previously in Eq. 4.

The proof of this lemma uses the following simple observation. Let Y be an arbitrary random variable taking values in \mathbb{N}_0 . Let X be a Rademacher random variable independent of Y . Then,

$$\mathbb{P}(Y \text{ has even parity}) = \mathbb{E}[X^Y], \quad \mathbb{P}(Y \text{ has odd parity}) = \mathbb{E}[X^{Y+1}|Y].$$

Appealing to this observation, we have,

$$\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]) = \mathbb{E} \left[\mathbf{X}_1^{\text{SUM}_\pi(\mathbf{C}, 1)} \dots \mathbf{X}_K^{\text{SUM}_\pi(\mathbf{C}, K)} \cdot \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right].$$

The expectation with respect to \mathbf{C} can now be evaluated by conditioning on $\|\mathbf{C}\|_1$ and using Proposition 7 along with the formula for the moment generating function for the Poisson and Multinomial distributions. A complete proof is provided in Appendix C.1. The advantage of the formula given in Lemma 8 is that it reduces the analysis of a combinatorial expression to controlling moments of subgaussian random variables. By expanding the exponential in Eq. 11 to a suitable order and using the fact that $\sqrt{d} \cdot \bar{X}_i$ are 1-subgaussian results in the following sharp upper bounds on the combinatorial coefficients $p_\pi(l_1, l_2 \dots l_K)$ and $\bar{p}_\pi(l_1, l_2 \dots l_K)$.

Proposition 9 *The coefficients $p_\pi(l_1, l_2 \dots l_K)$ satisfy the following estimates:*

$$p_\pi(l_1, l_2, \dots, l_K) \leq \begin{cases} C_k \cdot d^{-\frac{k}{2}} & : l_1 = l_2 = \dots = l_K = 0, \mathbf{o} = 0 \\ C_k \cdot d^{-k} & : l_1 = l_2 = \dots = l_K = 0, \mathbf{o} \geq 1 \\ C_k \cdot d^{-\frac{k + \sum_{i=1}^K l_i}{2}} & : l_i \leq |\pi^{-1}(i)|, l_i + |\pi^{-1}(i)| \text{ is even } \forall i \in [K] \\ C_k \cdot d^{-\frac{l_1 \vee l_2 \vee \dots \vee l_K}{2}} & : l_1 \vee l_2 \vee \dots \vee l_K \geq 2k \\ C_k \cdot d^{-k} & : \text{otherwise.} \end{cases}$$

For any $l_1, l_2, \dots, l_K \in \mathbb{N}_0$, we have, $\bar{p}_\pi(l_1, l_2, \dots, l_K) \leq p_\pi(l_1, l_2, \dots, l_K)$ and, in the special case $l_1 = l_2 = \dots = l_K = 0$, we have, $\bar{p}_\pi(0, 0, \dots, 0) \leq \frac{C_k}{d^k}$. In the above equations, C_k denotes a universal constant that depends only on k .

The proof of this Proposition can be found in Appendix C.2. Theorem 4 now follows immediately by substituting the above upper bounds on p_π and \bar{p}_π in Proposition 6 with $\epsilon \asymp n^{-\frac{1}{2}}$. More details can be found in Appendix D.

3. Optimal Statistical Query Procedures

Theorem 5 claims the existence of SQ algorithms which achieve the lower bounds proved in Theorem 4. A complete description of these procedures can be found in Appendix E. Interestingly, these procedures do not try and solve either the maximum likelihood problem or its spectral relaxation. Instead, these procedures leverage the ‘‘Partial Trace’’ technique used by Hopkins et al. (2016); Anandkumar et al. (2017); Biroli et al. (2019) to compress the empirical mean tensor in symmetric Tensor PCA into a vector or a matrix, and incur only a modest decrease in the signal-to-noise ratio. To implement these estimators in the SQ model, we leverage on the following result.

Fact 1 (SQ Scalar Mean Estimation; Feldman, 2017a) *There exists a SQ algorithm that given parameters $\xi, B > 0$, a query function $q : \bigotimes^k \mathbb{R}^d \mapsto \mathbb{R}$ and access to an arbitrary $\text{VSTAT}_D(n)$ oracle for any distribution $D \in \mathcal{D}(\pi) \cup \{D_0\}$ such that $\mathbb{E}_D q^2(\mathbf{T}) \leq B$, returns an estimate of $\mathbb{E}_D q(\mathbf{T})$ denoted by $\hat{\mathbb{E}}_D q(\mathbf{T})$ such that:*

$$\left| \hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T}) \right| \leq 8 \cdot \log(n) \cdot \sqrt{\frac{\text{Var}_D q(\mathbf{T})}{n}} + \xi.$$

after making at most $3 \log(4nB/\xi^2)$ queries.

We describe the optimal SQ algorithms in the following illustrative special cases.

Even order, Symmetric Tensor PCA Testing: In this case, $k = 2l$ for some $l \in \mathbb{N}$ and $K = 1$. The procedure uses Fact 1 to estimate the mean of the query

$$q(\mathbf{T}) = \mathbf{T}(\mathbf{I}_d, \mathbf{I}_d, \dots, \mathbf{I}_d) \stackrel{\text{def}}{=} \sum_{j_1, j_2, \dots, j_l \in [d]} T_{j_1, j_1, j_2, j_2, \dots, j_l, j_l}.$$

It is easy to check that when D is in the alternative hypothesis, $q(\mathbf{T}) \sim \mathcal{N}(1, \sqrt{d^k})$ and when $D = D_0$, then $q(\mathbf{T}) \sim \mathcal{N}(0, \sqrt{d^k})$. Consequently when $n \gg \sqrt{d^k}$, the error guarantee of the estimate from Fact 1 is small enough to ensure that null and alternate hypothesis are distinguishable.

Completely Asymmetric Tensor PCA: In this case $K = k$. The optimal procedure simply estimates the mean of each entry of the tensor using Fact 1. The resulting tensor of mean estimates is used as an estimate of the signal tensor for the estimation problem and its norm is used to solve the testing problem. Since the errors of a $\text{VSTAT}(n)$ oracle can be adversarial, the error in each entry can add up. Consequently, this approach works only when $n \gg d^k$.

Partially Symmetric Tensor PCA: The general case can be handled by a combination of the above techniques. At a high level, the signal tensor can be written as a outer product of a compressible part and an incompressible part. In the first step, a partial trace is used to compress the tensor by eliminating the compressible part. The mean of the incompressible

part is estimated entry-wise using Fact 1. The reduction in the size of the tensor during the compression phase ensures that the accumulation of adversarial SQ errors is minimized. The estimate of the incompressible part can be used to construct a test or to recover the compressible part of the signal tensor.

4. Effect of Noise Variance

In the formulation of Tensor PCA that we study in this paper, we assume that the noise variance is 1. It is natural to study formulations where the noise variance σ^2 is a parameter that is allowed to depend on n, d . More formally, in the testing problem $\text{TPCA}(n, d, \sigma^2)$, one is given a dataset consisting of n i.i.d. tensors $\mathbf{T}_{1:n} \in \bigotimes^k \mathbb{R}^d$, $k \geq 2$, and the goal is to determine if the tensors were drawn from the null hypothesis D_0 ,

$$D_0(\sigma^2) : (T_i)_{j_1, j_2, \dots, j_k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2) \quad \forall j_1, j_2, \dots, j_k \in [d], \forall i \in [n],$$

or from a distribution in the composite alternate hypothesis,

$$\mathcal{D}(\sigma^2) = \{D_{\mathbf{V}}(\sigma^2) : \mathbf{V} \in \bigotimes^k \mathbb{R}^d, \|\mathbf{V}\| = \sqrt{d^k}, \text{Rank}(\mathbf{V}) = 1\},$$

where under the alternate hypothesis $D_{\mathbf{V}}$ the entries of the tensor are distributed as

$$D_{\mathbf{V}}(\sigma^2) : (T_i)_{j_1, j_2, \dots, j_k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(\frac{V_{j_1, j_2, \dots, j_k}}{d^{\frac{k}{2}}}, \sigma^2\right) \quad \forall j_1, j_2, \dots, j_k \in [d], \forall i \in [n].$$

In the estimation problem $\text{TPCA}(n, d, \sigma^2)$, given a dataset $\mathbf{T}_{1:n}$ sampled i.i.d. from an unknown distribution $D \in \mathcal{D}(\sigma^2)$, the goal is to estimate the *signal tensor* $\mathbb{E}_D \mathbf{T} = \mathbf{V}/\sqrt{d^k}$. When allowed unrestricted access to the dataset $\text{TPCA}(n_1, d_1, \sigma_1^2)$ and $\text{TPCA}(n_2, d, \sigma_2^2)$ testing and estimation problems are statistically and computationally equivalent if $\sigma_1^2/n_1 = \sigma_2^2/n_2$ because of the following observation.

Observation 1 *Given n_1 samples drawn i.i.d. from $D_{\mathbf{V}}(\sigma_1^2)$ one can generate n_2 samples drawn i.i.d. from $D_{\mathbf{V}}(\sigma_2^2)$ without the knowledge of \mathbf{V} if $\sigma_1^2/n_1 = \sigma_2^2/n_2$. Moreover, this sampling mechanism is computationally efficient.*

Proof Note that one can make the transformation $(n_1, \sigma_1^2) \rightarrow (n_2, \sigma_2^2)$ in two steps: $(n_1, \sigma_1^2) \rightarrow (1, \sigma_1^2/n_1) = (1, \sigma_2^2/n_2) \rightarrow (n_2, \sigma_2^2)$. Hence, it is sufficient to show that given n samples from $D_{\mathbf{V}}(\sigma^2)$ one can generate 1 sample from $D_{\mathbf{V}}(\sigma^2/n)$ and vice-versa. Note that if $\mathbf{T}_{1:n} \stackrel{\text{i.i.d.}}{\sim} D_{\mathbf{V}}(\sigma^2)$ then,

$$\bar{T} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mathbf{T}_i \sim D_{\mathbf{V}}\left(\frac{\sigma^2}{n}\right).$$

Hence, to generate a sample from $D_{\mathbf{V}}(\sigma^2/n)$ one can compute the empirical average of the n samples from $D_{\mathbf{V}}(\sigma^2)$. For the other direction, we observe that since the empirical average is a sufficient statistic for the mean parameter \mathbf{V} of a Gaussian distribution (with known variance), the conditional law of $\mathbf{T}_{1:n}$ given \bar{T} does not depend on the mean parameter

\mathbf{V} . Hence, given one sample from $D_{\mathbf{V}}(\sigma^2/n)$ (i.e., $\overline{\mathbf{T}}$) one can generate n samples $\mathbf{T}_{1:n}$ by sampling from this conditional law without the knowledge of \mathbf{V} . Furthermore, since $\mathbf{T}_{1:n}, \overline{\mathbf{T}}$ are jointly Gaussian, this conditional law is Gaussian and can be efficiently sampled. ■

As a consequence of Observation 1, the problems $\text{TPCA}(n_1, d_1, \sigma_1^2)$ and $\text{TPCA}(n_2, d, \sigma_2^2)$ are statistically and computationally equivalent: given an algorithm \mathcal{A}_1 for $\text{TPCA}(n_1, d, \sigma_1^2)$ one can construct an algorithm \mathcal{A}_2 for $\text{TPCA}(n_2, d, \sigma_2^2)$ as follows: First use the reduction outlined in Observation 1 to construct a $\text{TPCA}(n_1, d, \sigma_1^2)$ instance from the $\text{TPCA}(n_2, d, \sigma_2^2)$ instance and then run \mathcal{A}_1 on the $\text{TPCA}(n_1, d, \sigma_1^2)$ instance. The procedure \mathcal{A}_2 has the same statistical performance on a $\text{TPCA}(n_2, d, \sigma_2^2)$ instance as procedure \mathcal{A}_1 has on $\text{TPCA}(n_1, d, \sigma_1^2)$ instance. Moreover if \mathcal{A}_1 is computationally efficient so is \mathcal{A}_2 .

Given the above observation, a natural question is whether the above equivalence is preserved when one can only access the data via an SQ oracle. Formally, let us consider a pair of distributions (D_1, D_2) such that:

$$D_1 \in \{D_0(\sigma_1^2)\} \cup \mathcal{D}(\sigma_1^2), D_2 \in \{D_0(\sigma_2^2)\} \cup \mathcal{D}(\sigma_2^2), \mathbb{E}_{D_1} \mathbf{T} = \mathbb{E}_{D_2} \mathbf{T}.$$

One can ask the question: Given access to a $\text{VSTAT}_{D_1}(n_1)$ can one simulate a $\text{VSTAT}_{D_2}(n_2)$ oracle and vice-versa? This would allow us to use SQ algorithms designed for $\text{TPCA}(n_2, d, \sigma_2^2)$ to solve $\text{TPCA}(n_1, d, \sigma_1^2)$ and vice-versa and lower bounds proved for one problem would have implications for the other problem.

Observation 2 *Suppose that $\sigma_2^2 = \sigma_1^2 \cdot S$ for some $S \in \mathbb{N}$. Then given access to a $\text{VSTAT}_{D_1}(n_1)$ oracle can one simulate a $\text{VSTAT}_{D_2}(n_2)$ oracle with:*

$$n_2 = \frac{Sn_1}{256 \log^2(n_1)} = \frac{1}{256 \log^2(n_1)} \cdot \frac{n_1 \sigma_2^2}{\sigma_1^2}.$$

Furthermore, this simulation makes at most $9 \log(n_1 S)$ queries to the $\text{VSTAT}_{D_1}(n_1)$ oracle to respond to one query for the $\text{VSTAT}_{D_2}(n_2)$ oracle.

The proof relies on a conditioning argument to reduce the variance of a statistic while preserving its expectation. This is often called Rao-Blackwellization (Rao, 1992; Blackwell, 1947) in statistics.

Proof Let $q : \otimes^k \mathbb{R}^d \rightarrow [0, 1]$ be the query for the $\text{VSTAT}_{D_2}(n_2)$ oracle. We will construct a query Q for the $\text{VSTAT}_{D_1}(n_1)$ oracle with the following properties:

1. Unbiasedness: $\mathbb{E}_{D_1} Q(\mathbf{T}) = \mathbb{E}_{D_2} q(\mathbf{T})$.
2. Variance bound: $\text{Var}_{D_1} Q(\mathbf{T}) \leq \text{Var}_{D_2} q(\mathbf{T})/S$.

Let us assume for the moment that such a query Q can be constructed and complete the proof of the claim. Indeed, by Fact 1 one can compute an estimate \hat{Q} by making at most $9 \log(n_1 S)$ queries to a $\text{VSTAT}_{D_1}(n_1)$ oracle which has the following error guarantee (we set

$\xi = (2n_2)^{-1}$:

$$\begin{aligned} \left| \hat{Q} - \mathbb{E}_{D_1} Q(\mathbf{T}) \right| &\leq 8 \log(n_1) \cdot \sqrt{\frac{\text{Var}_{D_1}(Q(\mathbf{T}))}{n_1}} + \xi \\ &\leq 2 \max \left(8 \log(n_1) \cdot \sqrt{\frac{\text{Var}_{D_1}(Q(\mathbf{T}))}{n_1}}, \xi \right) \\ &\leq \max \left(\sqrt{\frac{\text{Var}_{D_2}(q(\mathbf{T}))}{n_2}}, \frac{1}{n_2} \right). \end{aligned}$$

Since $\text{Var}_{D_2} q(\mathbf{T}) = \mathbb{E}_{D_2} q(\mathbf{T})^2 - (\mathbb{E}_{D_2} q(\mathbf{T}))^2 \leq (\mathbb{E}_{D_2} q(\mathbf{T})) \cdot (\mathbb{E}_{D_2} q(\mathbf{T}) - 1)$, \hat{Q} functions as a valid $\text{VSTAT}_{D_2}(n_2)$ response, completing the simulation of a $\text{VSTAT}_{D_2}(n_2)$ oracle.

In order to finish the proof, we need to describe the construction of the query Q with Properties 1 and 2 from above. In order to describe the construction, it will be helpful to realize a sample from D_1 and a sample from D_2 in the same probability space. We use the following natural construction: Let $\mathbf{T}_{1:S} \stackrel{\text{i.i.d.}}{\sim} D_2$ and define $\bar{\mathbf{T}} = \frac{1}{S} \sum_{i=1}^S \mathbf{T}_i$. Note that $\bar{\mathbf{T}} \sim D_1$. Note that, in order to construct the query Q , we need to find a function Q of $\bar{\mathbf{T}}$ with the properties:

1. Unbiasedness: $\mathbb{E}[Q(\bar{\mathbf{T}})] = \mathbb{E}[q(\mathbf{T}_1)]$,
2. Variance bound: $\text{Var}[Q(\bar{\mathbf{T}})] \leq \text{Var}[q(\mathbf{T}_1)]/S$.

Note that the function:

$$\bar{q}(\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_S) \stackrel{\text{def}}{=} \frac{1}{S} \sum_{i=1}^S q(\mathbf{T}_i).$$

satisfies Properties 1 and 2. However it is not a function of $\bar{\mathbf{T}}$. In order to fix this, define:

$$Q(\bar{\mathbf{T}}) \stackrel{\text{def}}{=} \mathbb{E} \left[\frac{1}{S} \sum_{i=1}^S q(\mathbf{T}_i) \middle| \bar{\mathbf{T}} \right].$$

This function Q satisfies Properties 1 and 2:

1. By the tower property:

$$\mathbb{E} Q(\bar{\mathbf{T}}) = \mathbb{E} \left[\frac{1}{S} \sum_{i=1}^S q(\mathbf{T}_i) \right] = \mathbb{E} q(\mathbf{T}_1).$$

2. Due to the fact that conditioning reduces variance (or Jensen's inequality):

$$\begin{aligned} \text{Var}[Q(\bar{\mathbf{T}})] &\leq \text{Var} \left[\frac{1}{S} \sum_{i=1}^S q(\mathbf{T}_i) \right] \\ &= \frac{\text{Var}[q(\mathbf{T}_1)]}{S}. \end{aligned}$$

Hence, we have constructed the desired function Q and this concludes the proof of the observation. \blacksquare

As a consequence of Observation 2, the SQ lower bounds proved for $\text{TPCA}(n, d, 1)$ in Theorem 4 can be lifted to lower bounds for $\text{TPCA}(n, d, \sigma^2)$ provided $\sigma^2 \geq \Omega_d(1)$.

Corollary 10 *Let $\sigma^2 \geq 1$. The following learning problems:*

1. $\text{TPCA}(n, d, \sigma^2)$ testing with $\circ = 0$ and sample size $n \ll \sigma^2 d^{\frac{k}{2}}$.
2. $\text{TPCA}(n, d, \sigma^2)$ estimation with $\circ = 0$ and sample size $n \ll \sigma^2 d^{\frac{k+2}{2}}$.
3. $\text{TPCA}(n, d, \sigma^2)$ testing or estimation with $\circ \geq 1$ and sample size $n \ll \sigma^2 d^{\frac{k+\circ}{2}}$.

cannot be solved in the SQ model with a query complexity which is polynomial in d . Furthermore, there exist SQ algorithms that make polynomial in d queries and solve:

1. *The Tensor PCA testing problem with $\circ = 0$ provided $n \gg \sigma^2 \sqrt{d^k}$.*
2. *The Tensor PCA estimation problem with $\circ = 0$ provided $n \gg \sigma^2 \sqrt{d^{k+2}}$.*
3. *The Tensor PCA testing and estimation problems with $\circ \geq 1$ provided $n \gg \sigma^2 \sqrt{d^{k+\circ}}$.*

Proof The lower bounds follow by a proof by contradiction: if there is a polynomial query SQ algorithm which beats the claimed lower bounds for $\text{TPCA}(n, d, \sigma^2)$, using the simulation argument in Observation 2, one can use it to solve $\text{TPCA}(n, d, 1)$ with a sample complexity n which beats the lower bounds in Theorem 4. For the upper bound, it is straightforward to extend the analysis of the optimal SQ procedures described in Section E to arbitrary noise level σ^2 . \blacksquare

Remark 11 *One can show that Corollary 10 holds for any $\sigma^2 \geq \Omega_d(1)$ (i.e., the constant 1 is not special). This is because the proof of Theorem 4 works for any constant σ^2 (not necessarily equal to 1).*

Observation 2 shows that an oracle with a larger sample size and a larger noise level is no stronger than an oracle with a smaller sample size and a proportionally smaller noise level. However the converse of this does not hold.

Observation 3 *There is a constant $c > 0$ small enough such that it is possible to solve symmetric $\text{TPCA}(n, d, \sigma^2)$ testing problem with $\sigma^2 = c/d$ in the SQ model with sample complexity $n = O_d(1)$ after making $O_d(1)$ queries.*

Proof Consider the query q :

$$q(\mathbf{T}) = \max_{\|\mathbf{x}\|=1} \sum_{i_1, i_2, \dots, i_k} T_{i_1, i_2, \dots, i_k} x_{i_1} x_{i_2} \dots x_{i_k}.$$

Lemma 2.1 in Richard and Montanari (2014) shows that if $\sigma^2 = c/d$ for a small enough constant c , then we have:

$$|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0(\sigma^2)} q(\mathbf{T})| \geq 0.5 \forall D \in \mathcal{D}(\sigma_2).$$

Furthermore Lemma 2.1 and 2.2 of Richard and Montanari show that:

$$\text{Var}_D q(\mathbf{T}) \leq \frac{C}{d}, \forall D \in \mathcal{D}(\sigma_2) \cup \{D_0(\sigma^2)\}.$$

for some constant $C > 0$. Consequently, by Fact 1 we can obtain an estimate $\hat{\mathbb{E}}_D q(\mathbf{T})$ such that $|\hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T})| \leq 0.1$ after making $O_d(1)$ queries to a $\text{VSTAT}_D(n)$ oracle with $n = O_d(1)$ and hence solve the testing problem. \blacksquare

The above observation shows that the assumption that $\sigma_2^2 \geq 1$ cannot be completely removed Corollary 10. Note that $n \geq \Omega_d(1)$ is information theoretic requirement on the sample size to solve $\text{TPCA}(n, d, \sigma^2)$ testing problem with $\sigma^2 = \Theta_d(1/d)$ and no known polynomial-time estimator achieves this sample complexity. Hence when $\sigma^2 = \Theta_d(1/d)$, SQ algorithms are so powerful that they completely fail to capture the conjectured computational-statistical gap and it is, in fact, possible to implement computationally intractable estimators in the SQ model.

The discussion in this section suggests that there is a possibility that for some well chosen setting of $1/d \ll \sigma^2 \ll 1$, the SQ lower bound may match the conjectured statistical-computational gap for all variants of Tensor PCA. We have not pursued this direction in our work because of the following reasons:

1. We believe that the setting $\sigma^2 = \Theta_d(1)$ is the most natural setting to study Tensor PCA in the SQ model. This is because Tensor PCA was intended to be a stylized setup to understand tensor based methods to fit latent variable models (Richard and Montanari, 2014). For example, consider estimating parameter $\boldsymbol{\mu}$ given data

$$\mathbf{x}_{1:n} \stackrel{\text{i.i.d.}}{\sim} 0.5\mathcal{N}(-\boldsymbol{\mu}, \sigma^2 \mathbf{I}_d) + 0.5\mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I}_d).$$

In these applications one observes several i.i.d. samples such that the signal-to-noise ratio per sample, measured by $\|\boldsymbol{\mu}\|^2/\sigma^2$ is $\Theta_d(1)$ and a sample size of $n \asymp d$ (the dimension of the parameter) is information theoretically necessary and sufficient. In Tensor PCA since the signal tensor has norm $\|\mathbb{E}\mathbf{T}\| = 1$, this suggests the natural setting of $\sigma^2 = \Theta_d(1)$ which is why we chose to study the $\text{TPCA}(n, d, 1)$ formulation. In contrast $\text{TPCA}(n, d, 1/d)$ seems very different from the situation arising in applications since the signal-to-noise ratio per sample diverges with d .

2. Setting σ^2 in a post-hoc manner to reproduce a widely believed computational-statistical gap does not seem like a satisfactory fix to the SQ model since this doesn't provide a framework to *predict* computational-statistical gaps in new problems.

5. Conclusions and Future Work

In this paper, we studied the Tensor PCA problem in the Statistical Query model. We found that SQ algorithms not only fail to solve Tensor PCA in the conjectured hard phase,

but they also have a strictly sub-optimal complexity compared to some polynomial time estimators such as the Richard-Montanari spectral estimator. Our analysis revealed that the optimal sample complexity in the SQ model depends on whether $\mathbb{E}\mathbf{T}$ is symmetric or not. Furthermore, we isolated a sample size regime in which only testing, but not estimation, is possible for the symmetric and even-order Tensor PCA problem in the SQ model.

Our results show that the statistical query framework leads to an incorrect and pessimistic prediction of the statistical-computation gap in Tensor PCA. An interesting problem for future work is to address this limitation of the statistical query framework. Feldman et al. (2018) found that the SQ framework leads to a similar pessimistic prediction of the statistical-computational gap in planted constraint satisfaction problems (CSPs). To address this, Feldman et al. proposed a hierarchical generalization of the $\text{VSTAT}(n)$ oracle for satisfiability problems known as the $\text{MVSTAT}(n, \ell)$ oracle where ℓ is a parameter controlling the strength of the oracle. The oracle is designed such that, for a well-chosen value of ℓ , it is possible to implement spectral estimators that achieve the best-known performance among known polynomial-time estimators. An interesting problem for future work is to identify a similar generalization of the $\text{VSTAT}(n)$ oracle for Tensor PCA. It would be exciting if this generalization is not tailored to a particular inference problem and works for a broad class of inference problems. Another desirable outcome is that this generalization is parameter-free, i.e., does not require setting parameters governing the strength of the oracle in a post-hoc manner to reproduce a widely believed computational-statistical gap. A modification to the SQ framework with these two properties would enable us to *predict* computational-statistical gaps in new problems rather than reproduce them. The recent work of Brennan et al. (2020) may provide hints towards identifying the correct modification to the SQ framework. This work shows that under some weak conditions, lower bounds for low-degree polynomial procedures (which are known to predict computational-statistical gaps correctly) are equivalent to a *sufficient criterion* for SQ lower bounds, namely the statistical dimension criterion of Feldman et al. (2017a). An interesting problem suggested by their work is to identify a relaxation of the current adversarial noise model used in the SQ framework under which the statistical dimension criterion of Feldman et al. (2017a) is a *necessary and sufficient* condition for SQ lower bounds. Such a modification to the SQ model would make SQ lower bounds equivalent to low-degree lower bounds, which are known to reliably predict the computational-statistical gap in a wide range of inference problems.

Acknowledgments

The authors would like to thank Tselil Schramm and Alex Wein for helpful discussions. DH was partially supported by NSF awards CCF-1740833, DMR-1534910, and IIS-1563785, a Google Faculty Research Award, and a Sloan Fellowship.

Appendix A. Framework for Statistical Query Lower Bounds

In this section, we review a few standard definitions and results used to prove SQ Lower bounds. First we review the notion of statistical dimension defined by Feldman et al. (2018).

Definition 12 (Statistical Dimension; Feldman et al., 2018) *Let $\epsilon > 0$ be an arbitrary tolerance parameter. Let D_\star be a fixed reference distribution and \mathcal{H} be a finite collection of distributions such that $D_\star \notin \mathcal{H}$. The statistical dimension at tolerance ϵ for D_\star v.s. \mathcal{H} denoted by $\text{SDN}(D_\star, \mathcal{H}, \epsilon)$ is the largest natural number m with the property: for any subset $\mathcal{S} \subset \mathcal{H}$ with size at least $|\mathcal{S}| \geq |\mathcal{H}|/m$, we have,*

$$\frac{1}{|\mathcal{S}|} \sum_{D \in \mathcal{S}} |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| \leq \epsilon \quad \forall q : \bigotimes_{k=1}^k \mathbb{R}^d \mapsto \mathbb{R}, \mathbb{E}_{D_\star} q^2(\mathbf{T}) = 1. \quad (12)$$

In our lower bounds, we will consider the following natural candidate for \mathcal{H} :

$$\mathcal{H}(\pi) \stackrel{\text{def}}{=} \{D_{\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \mathbf{v}_{\pi(k)}} : \mathbf{v}_{1:k} \in \{\pm 1\}^d\}.$$

Our lower bounds rely on the following result which shows that the statistical dimension is a lower bound on the query complexity of any SQ algorithm.

Proposition 13 (Feldman et al., 2018; Feldman, 2017b) *Consider the Tensor PCA problem with labelling function π . Let \mathcal{H} be any arbitrary finite subset of $\mathcal{D}(\pi)$. Any algorithm that solves the testing problem must make at least $\text{SDN}(D_0, \mathcal{H}(\pi), (3n)^{-\frac{1}{2}})$ queries. Furthermore, any SQ algorithm that solves the estimation problem must make at least $0.5 \cdot \text{SDN}(D_\star, \mathcal{H}(\pi), (3n)^{-\frac{1}{2}})$ queries, where D_\star is an arbitrary reference measure (not necessarily D_0).*

The claim of the above proposition for the testing problem is due to Feldman et al. (2018, Theorem 7.1). The claim regarding the estimation problem is due to Feldman (2017b, Theorem 4.2). A proof is provided in Appendix G for completeness.

Note in order to obtain lower bounds for the estimation problem we are free to choose D_\star . The natural choice of $D_\star = D_0$ will be sufficient to obtain tight lower bounds for most cases. For the remaining cases we will choose $D_\star = \overline{D}(\pi)$ which is a natural ‘average’ of all distributions in $\mathcal{H}(\pi)$.

$$\overline{D}(\pi) \stackrel{\text{def}}{=} D_{\overline{\mathbf{V}}(\pi)}, \quad \overline{\mathbf{V}}(\pi) = \mathbb{E}_{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\}^d)} [\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \mathbf{v}_{\pi(k)}].$$

The following lemma presents a lower bound on the statistical dimension. Similar results are implicit in various previous works on SQ lower bounds (Feldman et al., 2018).

Lemma 14 *For any $\epsilon > 0$ and any $D_\star \in \{D_0, \overline{D}(\pi)\}$,*

$$\text{SDN}(D_\star, \mathcal{H}(\pi), \epsilon) \geq \frac{\epsilon}{4} \cdot \left(\sup_{q: \mathbb{E}_{D_\star} q^2(\mathbf{T})=1} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| > \frac{\epsilon}{2} \right] \right)^{-1}.$$

The proof of the above lemma can also be found in Appendix G.

Appendix B. Proof of Proposition 6

As a consequence of Proposition 13 and Lemma 14 from the previous section, in order to prove a lower bound on the number of queries required to solve Tensor PCA with n samples, we need to prove an upper bound on the quantity:

$$\sup_{q: \mathbb{E}_{D_*} q^2(\mathbf{T})=1} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[\left| \mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_*} q(\mathbf{T}) \right| > \frac{1}{2\sqrt{3n}} \right].$$

This section is devoted to the proof of Proposition 6 which provides an analysis of the above quantity.

Proof Consider any query q such that $\mathbb{E}_{D_0} q^2(\mathbf{T}) = 1$. Observe that we have the following distributional equivalence:

$$D \sim \text{Unif}(\mathcal{H}(\pi)) \stackrel{d}{=} D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}, \mathbf{v}_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\}^d).$$

Hence,

$$\begin{aligned} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H})} [|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \epsilon] = \\ \mathbb{P}_{\mathbf{v}_{1:K} \sim \text{Unif}(\{\pm 1\}^d)} \left[\left| \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) \right| > \epsilon \right]. \end{aligned}$$

In the remainder of the proof, for the ease of notation, we will shorthand $\mathbb{E}_{\mathbf{v}_{1:K} \sim \text{Unif}(\{\pm 1\}^d)}$ and $\mathbb{P}_{\mathbf{v}_{1:K} \sim \text{Unif}(\{\pm 1\}^d)}$ simply as \mathbb{E} and \mathbb{P} .

Step 1, Markov Inequality: By Markov's Inequality, we have, for any $u \in \mathbb{N}$,

$$\mathbb{P} \left[\left| \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) \right| > \epsilon \right] \leq \frac{\mathbb{E} \left| \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) \right|^u}{\epsilon^u}. \quad (13)$$

Step 2, Fourier Expansion of $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$: Since the second moment of q is bounded, we can expand the query q in terms of the multivariate Hermite polynomials:

$$q(\mathbf{T}) = \sum_{\mathbf{c} \in \otimes^k \mathbb{N}_0^d} \hat{q}(\mathbf{c}) H_{\mathbf{c}}(\mathbf{T}).$$

We refer the reader to Appendix H.2 for definitions and basic properties of this orthonormal basis. By Parseval's identity, we have,

$$\mathbb{E}_{D_0} q^2(\mathbf{T}) = 1 \Leftrightarrow \sum_{\mathbf{c} \in \otimes^k \mathbb{N}_0^d} \hat{q}^2(\mathbf{c}) = 1.$$

The rational of expanding the query in this particular basis is the following elegant property (see Fact 5, Appendix H.2)²:

$$\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} H_{\mathbf{c}}(\mathbf{T}) = \prod_{j_1, j_2, \dots, j_k \in [d]} \left(\frac{(v_{\pi(1)})_{j_1} (v_{\pi(2)})_{j_2} \cdots (v_{\pi(k)})_{j_k}}{d^{\frac{k}{2}}} \right)^{c_{j_1, j_2, \dots, j_k}} \cdot \frac{1}{\sqrt{c_{j_1, j_2, \dots, j_k}!}}.$$

2. This property is also useful in the analysis of Tensor PCA via the low degree method (Hopkins et al., 2017; Hopkins, 2018; Kunisky et al., 2019), a heuristic technique motivated by Sum-of-Squares hierarchy to predict computational-statistical gaps.

After rearranging the terms in the product and recalling the $\text{SUM}(\cdot, \cdot)$, $\text{PARITY}(\cdot, \cdot)$ and entry-wise factorial and powering notations from Eqs. 5, 6 and 7 one obtains the following concise representation of the above formula:

$$\begin{aligned} \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} H_{\mathbf{c}}(\mathbf{T}) &= \frac{1}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}} \cdot \mathbf{v}_{\pi(1)}^{\text{SUM}(\mathbf{c}, 1)} \cdot \mathbf{v}_{\pi(2)}^{\text{SUM}(\mathbf{c}, 2)} \cdots \mathbf{v}_{\pi(k)}^{\text{SUM}(\mathbf{c}, k)} \\ &= \frac{1}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}} \cdot \mathbf{v}_1^{\text{SUM}_{\pi}(\mathbf{c}, 1)} \cdot \mathbf{v}_2^{\text{SUM}_{\pi}(\mathbf{c}, 2)} \cdots \mathbf{v}_K^{\text{SUM}_{\pi}(\mathbf{c}, K)} \\ &= \frac{1}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}} \cdot \mathbf{v}_1^{\text{PARITY}_{\pi}(\mathbf{c}, 1)} \cdot \mathbf{v}_2^{\text{PARITY}_{\pi}(\mathbf{c}, 2)} \cdots \mathbf{v}_K^{\text{PARITY}_{\pi}(\mathbf{c}, K)}. \end{aligned}$$

In the last step, we used the fact that $\mathbf{v}_i \in \{\pm 1\}^d$, $\forall i \in [K]$. Substituting this in the Fourier expansion of \mathbf{q} , one obtains the following representation of the bias $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} H_{\mathbf{c}}(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$ as a polynomial in the Boolean vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$:

$$\begin{aligned} \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} H_{\mathbf{c}}(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) &= \\ &= \sum_{\mathbf{c} \in \otimes^k \mathbb{N}_0^d; \|\mathbf{c}\|_1 \geq 1} \frac{\hat{q}(\mathbf{c})}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}} \cdot \mathbf{v}_1^{\text{PARITY}_{\pi}(\mathbf{c}, 1)} \cdots \mathbf{v}_K^{\text{PARITY}_{\pi}(\mathbf{c}, K)}. \end{aligned}$$

Observing that $\text{PARITY}(\mathbf{c}, l) \in \{0, 1\}^d$ for any \mathbf{c}, l we can regroup the terms in the above summation as follows:

$$\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) = \sum_{\mathbf{r}_1, \dots, \mathbf{r}_K \in \{0, 1\}^d} \tilde{q}_{\pi}(\mathbf{r}_1, \mathbf{r}_2 \cdots \mathbf{r}_K) \cdot \mathbf{v}_1^{\mathbf{r}_1} \cdot \mathbf{v}_2^{\mathbf{r}_2} \cdots \mathbf{v}_K^{\mathbf{r}_K}, \quad (14)$$

$$\tilde{q}_{\pi}(\mathbf{r}_1, \dots, \mathbf{r}_K) \stackrel{\text{def}}{=} \sum_{\substack{\mathbf{c} \in \otimes^k \mathbb{N}_0^d; \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_{\pi}(\mathbf{c}, i) = r_i \forall i \in [K]}} \frac{\hat{q}(\mathbf{c})}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}}. \quad (15)$$

The above representation will allow us to leverage tools from Fourier analysis over the Boolean hypercube in the following step.

Step 3, Hypercontractivity: We define the following polynomial in $\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K \in \{\pm 1\}^d$, for any $u \in \mathbb{N}$, $u \geq 2$.

$$P_u(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) \stackrel{\text{def}}{=} \sum_{\mathbf{r}_1, \dots, \mathbf{r}_K \in \{0, 1\}^d} (\sqrt{u-1})^{\|\mathbf{r}_1\|_1 + \dots + \|\mathbf{r}_K\|_1} \cdot \tilde{q}_{\pi}(\mathbf{r}_{1:K}) \cdot \mathbf{v}_1^{\mathbf{r}_1} \cdot \mathbf{v}_2^{\mathbf{r}_2} \cdots \mathbf{v}_K^{\mathbf{r}_K},$$

where the coefficients \tilde{q} have been defined in Eq. 15. Note that the bias $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$ is a smoothed out version of the polynomial P (cf. Eq. 14) in the following sense:

$$\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) = \text{SMOOTH}_{\sqrt{u-1}}[P](\mathbf{v}_1, \mathbf{v}_2 \cdots, \mathbf{v}_K),$$

where, for any $\lambda \geq 1$, the operator $\text{SMOOTH}_{\lambda}[\cdot]$ down weights the coefficient of $\mathbf{v}_1^{\mathbf{r}_1} \cdots \mathbf{v}_K^{\mathbf{r}_K}$ by $\lambda^{-(\|\mathbf{r}_1\|_1 + \dots + \|\mathbf{r}_K\|_1)}$ (see Appendix H.1 for more background). $(2, u)$ -Hypercontractive inequalities allow us to upper bound the L_u norm of $\text{SMOOTH}_{\sqrt{u-1}}[P](\mathbf{v}_1, \mathbf{v}_2 \cdots, \mathbf{v}_K)$ in

terms of the L_2 norm of $P(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K)$. Applying the $(2, u)$ -Hypercontractive inequality (Fact 4), we obtain,

$$\mathbb{E}|\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} - \mathbb{E}_{D_0} q(\mathbf{T})|^u \leq (\mathbb{E}P_u^2(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K))^{\frac{u}{2}}. \quad (16)$$

Note that the monomial functions $\mathbf{v}_1^{r_1} \mathbf{v}_2^{r_2} \cdots \mathbf{v}_K^{r_K}$ are orthonormal when $\mathbf{v}_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\}^d)$ (See again Appendix H.1). Hence, by Parseval's identity and Eq. 15,

$$\begin{aligned} \mathbb{E}P_u^2(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) &= \sum_{\mathbf{r}_1, \dots, \mathbf{r}_K \in \{0,1\}^d} (u-1)^{\|\mathbf{r}_1\|_1 + \|\mathbf{r}_2\|_1 + \dots + \|\mathbf{r}_K\|_1} \cdot \hat{q}_\pi^2(\mathbf{r}_1, \mathbf{r}_2 \cdots \mathbf{r}_K) \\ &= \sum_{\mathbf{r}_1, \dots, \mathbf{r}_K \in \{0,1\}^d} (u-1)^{\|\mathbf{r}_1\|_1 + \|\mathbf{r}_2\|_1 + \dots + \|\mathbf{r}_K\|_1} \cdot \left(\sum_{\substack{\mathbf{c} \in \otimes^k \mathbb{N}_0^d: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = \mathbf{r}_i \ \forall i \in [K]}} \frac{\hat{q}(\mathbf{c})}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}} \right)^2. \end{aligned}$$

Applying the Cauchy-Schwarz inequality,

$$\begin{aligned} \mathbb{E}P_u^2(\mathbf{v}_{1:K}) &\leq \\ &\sum_{\mathbf{r}_1, \dots, \mathbf{r}_K} (u-1)^{\|\mathbf{r}_1\|_1 + \dots + \|\mathbf{r}_K\|_1} \cdot \left(\sum_{\substack{\mathbf{c}: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = \mathbf{r}_i \ \forall i}} \frac{1}{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!} \right) \cdot \left(\sum_{\substack{\mathbf{c}: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = \mathbf{r}_i \ \forall i}} \hat{q}^2(\mathbf{c}) \right). \end{aligned}$$

Note that, since $\mathbb{E}_{D_0} q^2(\mathbf{T}) = 1$,

$$\sum_{\mathbf{r}_1, \dots, \mathbf{r}_K} \sum_{\substack{\mathbf{c}: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = \mathbf{r}_i \ \forall i}} \hat{q}^2(\mathbf{c}) \leq 1.$$

Hence by the ℓ_1/ℓ_∞ Hölder's inequality:

$$\mathbb{E}P_u^2(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) \leq \max_{\mathbf{r}_1, \mathbf{r}_2 \cdots \mathbf{r}_K \in \{0,1\}^d} (u-1)^{\|\mathbf{r}_1\|_1 + \dots + \|\mathbf{r}_K\|_1} \cdot \left(\sum_{\substack{\mathbf{c}: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = \mathbf{r}_i \ \forall i}} \frac{1}{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!} \right).$$

By a relabelling of $[d]$, the expression inside the maximum in the above display depends on $\mathbf{r}_1, \mathbf{r}_2 \cdots \mathbf{r}_K$ only through $l_1 \stackrel{\text{def}}{=} \|\mathbf{r}_1\|_1, l_2 \stackrel{\text{def}}{=} \|\mathbf{r}_2\|_1 \cdots, l_K \stackrel{\text{def}}{=} \|\mathbf{r}_K\|_1$. Hence we have,

$$\mathbb{E}P_u^2(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) \leq \max_{l_1, l_2 \cdots l_K \in \mathbb{N}_0} \left(e \cdot (u-1)^{l_1 + l_2 \cdots l_K} \cdot p_\pi(l_1, l_2 \cdots l_K) \right),$$

where,

$$p_\pi(l_1, l_2 \cdots l_K) \stackrel{\text{def}}{=} \frac{1}{e} \sum_{\substack{\mathbf{c}: \|\mathbf{c}\|_1 \geq 1 \\ \text{PARITY}_\pi(\mathbf{c}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \ \forall i \in [K]}} \frac{1}{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}.$$

Substituting the above bound on the second moment in Eq. 16, we obtain,

$$\mathbb{E}|\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})|^u \leq \left(\max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left(e \cdot (u-1)^{l_1 + l_2 \dots l_K} \cdot p_\pi(l_1, l_2 \dots l_K) \right) \right)^{\frac{u}{2}}.$$

Finally Markov's inequality (cf. Eq. 13) gives us,

$$\mathbb{P} \left[\left| \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T}) \right| > \epsilon \right] \leq \left(\frac{1}{\epsilon^2} \cdot \max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left(e \cdot (u-1)^{l_1 + l_2 \dots l_K} \cdot p_\pi(l_1, l_2 \dots l_K) \right) \right)^{\frac{u}{2}}.$$

This concludes the proof of the proposition for the reference measure D_0 . The result for the reference measure $\bar{D}(\pi)$ follows via a similar calculation and we only sketch the argument for this case. Note that if $\mathbf{T} \sim D_0$, then,

$$\mathbf{T} + \frac{\bar{\mathbf{V}}(\pi)}{\sqrt{d^k}} \sim \bar{D}(\pi).$$

Now consider any query function such that $\mathbb{E}_{\bar{D}(\pi)} q^2(\mathbf{T}) = 1$. This means that:

$$\mathbb{E}_{D_0} q^2 \left(\mathbf{T} + \frac{\bar{\mathbf{V}}(\pi)}{\sqrt{d^k}} \right) = 1.$$

Hence we can expand the function $q\left(\frac{\bar{\mathbf{V}}(\pi)}{\sqrt{d^k}} + \cdot\right)$ in the multivariate Hermite Basis:

$$q \left(\frac{\bar{\mathbf{V}}(\pi)}{\sqrt{d^k}} + \mathbf{T} \right) = \sum_{\mathbf{c}} \hat{q}(\mathbf{c}) H_{\mathbf{c}}(\mathbf{T}).$$

Hence,

$$\mathbb{E}_{\bar{D}(\pi)} q(\mathbf{T}) = \hat{q}(\mathbf{0}), \quad \mathbb{E}_{\bar{D}(\pi)} q^2(\mathbf{T}) = \sum_{\mathbf{c}} \hat{q}(\mathbf{c})^2 = 1.$$

Define the tensor $\mathbf{V}(\pi)$ as:

$$\mathbf{V}(\pi) = \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)}.$$

Also note that when $\mathbf{T} \sim D_0$, then,

$$\mathbf{T} + \frac{\mathbf{V}(\pi)}{\sqrt{d^k}} \sim D_{\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)}}.$$

Hence,

$$\begin{aligned} \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) &= \mathbb{E}_{D_0} q \left(\mathbf{T} + \frac{\mathbf{V}(\pi)}{\sqrt{d^k}} \right) \\ &= \mathbb{E}_{D_0} \left[\sum_{\mathbf{c}} \hat{q}(\mathbf{c}) H_{\mathbf{c}} \left(\mathbf{T} + \frac{\mathbf{V}(\pi) - \bar{\mathbf{V}}(\pi)}{\sqrt{d^k}} \right) \right]. \end{aligned}$$

Due to the fact that each entry of $\mathbf{V}(\pi)$ is either 1 or marginally Rademacher distributed and $\mathbb{E}\mathbf{V}(\pi) = \overline{\mathbf{V}}(\pi)$, we have,

$$V(\pi)_{i_1, i_2, \dots, i_k} - \overline{V}(\pi)_{i_1, i_2, \dots, i_k} = \begin{cases} 0 & : (i_1, i_2 \dots i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi)) \\ V(\pi)_{i_1, i_2, \dots, i_k} & : (i_1, i_2 \dots i_k) \notin \text{Supp}(\overline{\mathbf{V}}(\pi)) \end{cases}.$$

Using Fact 5, we can compute,

$$\mathbb{E}_{D_0} H_{\mathbf{c}} \left(\mathbf{T} + \frac{\mathbf{V}(\pi) - \overline{\mathbf{V}}(\pi)}{\sqrt{d^k}} \right) = \begin{cases} 0 & : \text{Supp}(\mathbf{c}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) \neq \phi \\ \frac{\mathbf{V}(\pi)^{\mathbf{c}}}{\sqrt{c! d^{k\|\mathbf{c}\|_1}}} & : \text{Supp}(\mathbf{c}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi \end{cases}.$$

Hence we obtain the following Fourier expansion of the bias $\mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \dots \otimes \mathbf{v}_{\pi(k)}}} q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})$:

$$\begin{aligned} \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \dots \otimes \mathbf{v}_{\pi(k)}}} H_{\mathbf{c}}(\mathbf{T}) - \mathbb{E}_{\overline{D}(\pi)} q(\mathbf{T}) = \\ \sum_{\substack{\mathbf{c} \in \otimes^k \mathbb{N}_0^d, \|\mathbf{c}\|_1 \geq 1 \\ \text{Supp}(\mathbf{c}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi}} \frac{\hat{q}(\mathbf{c}) \cdot \mathbf{v}_1^{\text{PARITY}_{\pi}(\mathbf{c}, 1)} \dots \mathbf{v}_K^{\text{PARITY}_{\pi}(\mathbf{c}, K)}}{\sqrt{d^{k\|\mathbf{c}\|_1} \cdot \mathbf{c}!}}. \end{aligned}$$

From here on wards the arguments made previously for the case when the reference distribution was D_0 can be repeated verbatim and give the second claim of the proposition. \blacksquare

Appendix C. Asymptotic Analysis of p_{π} and \overline{p}_{π}

Recall that Proposition 6 showed that the number of queries required to solve the Tensor PCA problems was intimately connected to the asymptotics of the combinatorial coefficients $p_{\pi}(l_1, l_2, \dots, l_K)$ and $\overline{p}_{\pi}(l_1, l_2, \dots, l_K)$. We had also arrived at a probabilistic interpretation of these coefficients as:

$$p_{\pi}(l_1, l_2, \dots, l_K) = \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_{\pi}(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]),$$

and,

$$\begin{aligned} \overline{p}_{\pi}(l_1, l_2, \dots, l_K) \\ = \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_{\pi}(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi), \end{aligned}$$

where $\mathbf{C} \in \otimes^k \mathbb{N}_0^d$ was a tensor with i.i.d. $\text{Pois}(d^{-k})$ entries. This goal of this section is to prove Proposition 9 which provides an upper bound on these combinatorial coefficients. This section is organized as follows:

- Section C.1 contains the proof of Lemma C.1 which provides a convenient analytic formula for the probabilities $\mathbb{P}(\text{PARITY}_{\pi}(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K])$.
- Section C.2 uses this formula to prove Proposition 9.

C.1 Proof of Lemma 8

Proof The proof of this lemma relies on the following simple observation. Let Y be an arbitrary random variable taking values in \mathbb{N}_0 . Let X be a Rademacher random variable independent of Y . Then,

$$\mathbb{I}(Y \text{ has even parity}) = \mathbb{E}[X^Y | Y], \quad \mathbb{I}(Y \text{ has odd parity}) = \mathbb{E}[X^{Y+1} | Y].$$

Appealing to this observation, we have,

$$\mathbb{I}(\text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]) = \mathbb{E} \left[\mathbf{X}_1^{\text{SUM}_\pi(\mathbf{C}, 1)} \cdots \mathbf{X}_K^{\text{SUM}_\pi(\mathbf{C}, K)} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \mid \mathbf{C} \right].$$

Recalling Eq. 5, the above expression can be written as:

$$\mathbb{I}(\text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]) = \mathbb{E} \left[\mathbf{X}_{\pi(1)}^{\text{SUM}(\mathbf{C}, 1)} \cdots \mathbf{X}_{\pi(k)}^{\text{SUM}(\mathbf{C}, k)} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \mid \mathbf{C} \right].$$

Hence,

$$\begin{aligned} \mathbb{P}(\text{PARITY}(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]) &= \mathbb{E} \left[\mathbf{X}_{\pi(1)}^{\text{SUM}(\mathbf{C}, 1)} \cdot \mathbf{X}_{\pi(2)}^{\text{SUM}(\mathbf{C}, 2)} \cdots \mathbf{X}_{\pi(k)}^{\text{SUM}(\mathbf{C}, k)} \cdot \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right] \\ &= \mathbb{E} \mathbb{E} \left[\mathbf{X}_{\pi(1)}^{\text{SUM}(\mathbf{C}, 1)} \cdot \mathbf{X}_{\pi(2)}^{\text{SUM}(\mathbf{C}, 2)} \cdots \mathbf{X}_{\pi(k)}^{\text{SUM}(\mathbf{C}, k)} \cdot \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \mid \|\mathbf{C}\|_1, \mathbf{X}_1, \mathbf{X}_2 \cdots \mathbf{X}_K \right] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\left(\prod_{i=1}^K \mathbb{E} \left[\mathbf{X}_{\pi(i)}^{\text{SUM}(\mathbf{C}, i)} \mid \|\mathbf{C}\|_1, \mathbf{X}_{\pi(i)} \right] \right) \cdot \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right] \\ &\stackrel{(b)}{=} \mathbb{E} \left[(\bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \cdots \bar{X}_K^{s_K}) \|\mathbf{C}\|_1 \cdot \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right] \\ &\stackrel{(c)}{=} \frac{1}{e} \cdot \mathbb{E} \left[e^{\bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \cdots \bar{X}_K^{s_K}} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right]. \end{aligned}$$

In the above display, the Equation marked (a) uses the conditional independence implied by Proposition 7. In the step marked (b), we used the formula for the generating function of a Multinomial distribution (Fact 2). The equality (c) uses the formula for the generating function of a Poisson distribution (Fact 3). This proves the claim of the lemma. \blacksquare

C.2 Proof of Proposition 9

This section is devoted to the proof of Proposition 9. Recall that Lemma 8 relates p_π, \bar{p}_π to expectations of the form:

$$\mathbb{E} \left[e^{\bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \cdots \bar{X}_K^{s_K} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij}} \right].$$

Our strategy to analyze the above expectation will to approximate the exponential by its Taylor expansion of a suitable degree. Recalling the independence of the Rademacher vectors $\mathbf{X}_{1:K}$, this reduces our task to analyze expectations of the form:

$$\mathbb{E} \left[\bar{X}^s \prod_{i=1}^l X_i \right].$$

This is the content of the following intermediate lemma.

Lemma 15 *Let $\mathbf{X} \in \{\pm 1\}^d$ be a random vector with i.i.d. Rademacher entries. Let \bar{X} be the average of the entries of \mathbf{X} . Then, if $l > s$ or if $l + s$ has odd parity, then,*

$$\mathbb{E} \left[\bar{X}^s \prod_{i=1}^l X_i \right] = 0.$$

Otherwise,

$$0 \leq \mathbb{E} \left[\bar{X}^s \prod_{i=1}^l X_i \right] \leq \frac{C_s}{\sqrt{d^{l+s}}},$$

where C_s denotes a constant depending only on s .

Proof When $l + s$ has odd parity, then $\bar{X}^s \prod_{i=1}^l X_i$ is an odd polynomial in \mathbf{X} and hence,

$$\mathbb{E}[\bar{X}^s \prod_{i=1}^l X_i] = 0, \text{ if } l + s \text{ is odd.}$$

When $l > s$, note that \bar{X}^s is a degree s polynomial, and since $s < l$, it is orthogonal to $\prod_{i=1}^l X_i$, hence,

$$\mathbb{E}[\bar{X}^s \prod_{i=1}^l X_i] = 0, \text{ if } l > s.$$

Finally when $l \leq s$, observe that $\mathbb{E}[\bar{X}^s \prod_{i=1}^l X_i]$ is simply the coefficient of $\prod_{i=1}^l X_i$ in the expansion of \bar{X}^s in the monomial basis for Boolean functions (see Appendix H.1). Expanding \bar{X}^s using the Multinomial theorem shows that it is a polynomial with all non-negative coefficients. This gives us the claim:

$$\mathbb{E}[\bar{X}^s \prod_{i=1}^l X_i] \geq 0.$$

In order to understand its asymptotics, we decompose \bar{X} as:

$$\bar{X} = \frac{Z}{d} + \frac{d-l}{d}\bar{Y}, \quad Z \stackrel{\text{def}}{=} X_1 + X_2 + \cdots + X_l, \quad \bar{Y} = \frac{\sum_{i=l+1}^d X_i}{d-l}.$$

Hence,

$$\mathbb{E}\bar{X}^s \prod_{i=1}^l X_i = \sum_{t=0}^s \binom{s}{t} \cdot \left(\frac{d-l}{d}\right)^{s-t} \cdot \frac{1}{d^t} \cdot \mathbb{E} \left[Z^t \prod_{i=1}^l X_i \right] \cdot \mathbb{E}\bar{Y}^{s-l}.$$

We observe that (using the same argument made previously), when $t < l$, Z^t is a degree t polynomial in X_1, X_2, \dots, X_l and hence orthogonal to $\prod_{i=1}^l X_i$:

$$\mathbb{E} \left[Z^t \prod_{i=1}^l X_i \right] = 0.$$

Hence,

$$\mathbb{E}\bar{X}^s \prod_{i=1}^l X_i = \sum_{t=l}^s \binom{s}{t} \cdot \left(\frac{d-l}{d}\right)^{s-t} \cdot \frac{1}{d^t} \cdot \mathbb{E} \left[Z^t \prod_{i=1}^l X_i \right] \cdot \mathbb{E}\bar{Y}^{s-t} \quad (17)$$

Next we observe that,

$$|Z^t \prod_{i=1}^l X_i| \leq l^t,$$

Since $\sqrt{d-l}\bar{Y}$ is 1-subgaussian (since it is a normalized sum of i.i.d. Rademacher variables), we have,

$$\mathbb{E}|\bar{Y}^{s-t}| \leq \frac{1}{\sqrt{(d-l)^{s-t}}} \cdot 2^{s-t} \cdot (s-t)^{\frac{s-t}{2}} \leq \frac{C_s}{\sqrt{d^{s-t}}},$$

where C_s denotes a constant depending only on s . Substituting these estimates in the Eq. 17, we obtain,

$$\mathbb{E}\bar{X}^s \prod_{i=1}^l X_i \leq \frac{C_s}{\sqrt{d^{s+t}}},$$

for some constant C_s depending only on s . This concludes the proof of the lemma. \blacksquare

We now present the proof of Proposition 9 restated below for convenience.

Proposition 16 *The coefficients $p_\pi(l_1, l_2, \dots, l_K)$ satisfy the following estimates:*

$$p_\pi(l_1, l_2, \dots, l_K) \leq \begin{cases} C_k \cdot d^{-\frac{k}{2}} & : l_1 = l_2 = \cdots = l_K = 0, \mathfrak{o} = 0 \\ C_k \cdot d^{-k} & : l_1 = l_2 = \cdots = l_K = 0, \mathfrak{o} \geq 1 \\ C_k \cdot d^{-\frac{k + \sum_{i=1}^K l_i}{2}} & : l_i \leq |\pi^{-1}(i)|, l_i + |\pi^{-1}(i)| \text{ is even } \forall i \in [K] \\ C_k \cdot d^{-\frac{l_1 \vee l_2 \vee \cdots \vee l_K}{2}} & : l_1 \vee l_2 \vee \cdots \vee l_K \geq 2k \\ C_k \cdot d^{-k} & : \text{otherwise.} \end{cases}$$

For any $l_1, l_2, \dots, l_K \in \mathbb{N}_0$, we have, $\bar{p}_\pi(l_1, l_2, \dots, l_K) \leq p_\pi(l_1, l_2, \dots, l_K)$ and, in the special case $l_1 = l_2 = \dots = l_K = 0$, we have, $\bar{p}_\pi(0, 0, \dots, 0) \leq \frac{C_k}{d^k}$. In the above equations, C_k denotes a universal constant that depends only on k .

Proof First recall the definitions of the coefficients p_π and \bar{p}_π :

$$\begin{aligned} p_\pi(l_{1:K}) &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]), \\ \bar{p}_\pi(l_{1:K}) &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\bar{\mathbf{V}}(\pi)) = \emptyset). \end{aligned}$$

The above expressions immediately give us $\bar{p}_\pi(l_1, l_2 \dots l_K) \leq p_\pi(l_1, l_2 \dots l_K)$. Next we split our analysis into 4 cases:

Case 1: $l_1 \vee l_2 \vee \dots \vee l_K = 0$. In this case, Lemma 8 gives us:

$$\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) = \frac{1}{e} \mathbb{E} \left[e^{\bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \dots \bar{X}_K^{s_K}} \right],$$

where s_i were defined as:

$$s_i \stackrel{\text{def}}{=} |\pi^{-1}(i)| = |\{j \in [k] : \pi(j) = i\}|.$$

Using the inequality $|e^x - 1 - x| \leq \frac{ex^2}{2} \forall x \leq 1$ one obtains:

$$\begin{aligned} &\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) \\ &\leq \frac{1}{e} \left(1 + \mathbb{E} \bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \dots \bar{X}_K^{s_K} + \frac{e \cdot \mathbb{E} \bar{X}_1^{2s_1} \cdot \bar{X}_2^{2s_2} \dots \bar{X}_K^{2s_K}}{2} \right) \\ &= \frac{1}{e} + \frac{1}{2 \cdot d^k}. \end{aligned}$$

By Lemma 15,

$$\mathbb{E} \bar{X}_1^{2s_1} \cdot \bar{X}_2^{2s_2} \dots \bar{X}_K^{2s_K} \leq \frac{C_k}{d^{s_1+s_2+\dots+s_K}} \leq \frac{C_k}{d^k}.$$

Recall that \circ denoted the number of $i \in [K]$ such that s_i is odd. If $\circ \geq 1$, then,

$$\mathbb{E} \bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \dots \bar{X}_K^{s_K} = 0.$$

Otherwise, if $\circ = 0$, then Lemma 15 gives us,

$$|\mathbb{E} \bar{X}_1^{s_1} \cdot \bar{X}_2^{s_2} \dots \bar{X}_K^{s_K}| \leq \frac{C_k}{\sqrt{d^k}}.$$

Hence,

$$\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) \leq \begin{cases} \frac{1}{e} + \frac{C_k}{\sqrt{d^k}} & : \circ = 0 \\ \frac{1}{e} + \frac{C_k}{d^k} & : \circ \geq 1 \end{cases}.$$

We use the above estimates to bound $p_\pi(0, 0, \dots, 0)$ and $\bar{p}_\pi(0, 0, \dots, 0)$. Recall that,

$$\begin{aligned} p_\pi(0, 0, \dots, 0) &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) \\ &= \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) - \mathbb{P}(\|\mathbf{C}\|_1 = 0) \\ &\stackrel{(a)}{=} \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) - \frac{1}{e} \\ &\leq \begin{cases} \frac{C_k}{\sqrt{d^k}} & : \circ = 0 \\ \frac{C_k}{d^k} & : \circ \geq 1 \end{cases}. \end{aligned}$$

In the equation marked (a) we used the fact that $\|\mathbf{C}\|_1 \sim \text{Pois}(1)$ (cf. Proposition 7). Next we analyze $\bar{p}_\pi(0, 0, \dots, 0)$. We have,

$$\begin{aligned} \bar{p}_\pi(0, \dots, 0) &= \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi) \\ &= \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi) - \mathbb{P}(\|\mathbf{C}\|_1 = 0) \\ &= \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K], \text{Supp}(\mathbf{C}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi) - \frac{1}{e}. \end{aligned}$$

Note that,

$$\text{Supp}(\mathbf{C}) \cap \text{Supp}(\overline{\mathbf{V}}(\pi)) = \phi \Leftrightarrow \sum_{(i_1, i_2, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} = 0.$$

Furthermore, the random variables,

$$\sum_{(i_1, i_2, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} \text{ and } \text{PARITY}_\pi(\mathbf{C}, \cdot),$$

are independent since any C_{i_1, i_2, \dots, i_k} with $(i_1, i_2, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))$ is counted an even number of times in $\text{SUM}_\pi(\mathbf{C}, \cdot)$ and hence does not influence the parities $\text{PARITY}_\pi(\mathbf{C}, \cdot)$. Hence,

$$\begin{aligned} \bar{p}_\pi(0, 0, \dots, 0) &= \\ &\mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) \cdot \mathbb{P}\left(\sum_{(i_1, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} = 0\right) - \frac{1}{e}. \end{aligned}$$

Furthermore,

$$\sum_{(i_1, i_2, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} \sim \text{Pois}\left(\frac{|\text{Supp}(\overline{\mathbf{V}}(\pi))|}{d^k}\right),$$

and, since the entries of $\overline{\mathbf{V}}(\pi)$ are in $\{0, 1\}$:

$$\begin{aligned} \frac{1}{d^k} \cdot |\text{Supp}(\overline{\mathbf{V}}(\pi))| &= \frac{1}{d^k} \sum_{i_1, i_2, \dots, i_k} \overline{\mathbf{V}}(\pi)_{i_1, i_2, \dots, i_k} \\ &= \frac{1}{d^k} \sum_{i_1, i_2, \dots, i_k} (\mathbb{E}_{\mathbf{X}_{1:K} \sim \text{Unif}(\{\pm 1\}^d)} \mathbf{X}_{\pi(1)} \otimes \mathbf{X}_{\pi(2)} \cdots \otimes \mathbf{X}_{\pi(k)})_{i_1, i_2, \dots, i_k} \\ &= \mathbb{E} \overline{X}_1^{s_1} \cdot \overline{X}_2^{s_2} \cdots \overline{X}_K^{s_K}. \end{aligned}$$

Hence,

$$\mathbb{P} \left(\sum_{(i_1, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} = 0 \right) = \exp(-\mathbb{E} \overline{X}_1^{s_1} \cdot \overline{X}_2^{s_2} \cdots \overline{X}_K^{s_K}).$$

This gives us,

$$\begin{aligned} &\overline{p}_\pi(0, 0, \dots, 0) \\ &= \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = \mathbf{0} \forall i \in [K]) \cdot \mathbb{P} \left(\sum_{(i_1, \dots, i_k) \in \text{Supp}(\overline{\mathbf{V}}(\pi))} C_{i_1, i_2, \dots, i_k} = 0 \right) - \frac{1}{e} \\ &= \frac{1}{e} (\mathbb{E} \exp(\overline{X}_1^{s_1} \cdot \overline{X}_2^{s_2} \cdots \overline{X}_K^{s_K}) - \mathbb{E} \overline{X}_1^{s_1} \cdot \overline{X}_2^{s_2} \cdots \overline{X}_K^{s_K}) \\ &\leq \frac{1}{e} \cdot \left(\mathbb{E}(\overline{X}_1^{s_1} \cdots \overline{X}_K^{s_K}) - \mathbb{E} \overline{X}_1^{s_1} \cdots \mathbb{E} \overline{X}_K^{s_K} + \frac{e \mathbb{E}(\overline{X}_1^{s_1} \cdots \overline{X}_K^{s_K} - \mathbb{E} \overline{X}_1^{s_1} \cdots \mathbb{E} \overline{X}_K^{s_K})^2}{2} \right) \\ &\leq \frac{C_k}{d^k} \end{aligned}$$

In the last step, we again appealed to Lemma 15. Hence,

$$\overline{p}_\pi(0, 0, \dots, 0) \leq \frac{C_k}{d^k}.$$

This concludes the analysis of this case.

Case 2: $l_i \leq s_i$ $i = 1, 2, \dots, K$, $l_1 \vee l_2 \cdots \vee l_K \geq 1$. Recall that,

$$p_\pi(l_1, l_2, \dots, l_K) = \mathbb{P}(\|\mathbf{C}\|_1 \geq 1, \text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]).$$

Since $l_1 \vee l_2 \cdots \vee l_K \geq 1$, can remove the condition $\|\mathbf{C}\|_1 \geq 1$:

$$p_\pi(l_1, l_2, \dots, l_K) = \mathbb{P}(\text{PARITY}_\pi(\mathbf{C}, i) = (\mathbf{1}_{l_i}, \mathbf{0}) \forall i \in [K]).$$

By Lemma 8, we have,

$$p_\pi(l_1, l_2, \dots, l_K) = \frac{1}{e} \mathbb{E} \left[e^{\overline{X}_1^{s_1} \cdot \overline{X}_2^{s_2} \cdots \overline{X}_K^{s_K}} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right].$$

Using the bound:

$$1 + x \leq e^x \leq 1 + x + \frac{ex^2}{2}, \quad \forall x \leq 1,$$

we obtain,

$$p_\pi(l_1, \dots, l_K) \leq \frac{1}{e} \cdot \left(\prod_{i=1}^K \prod_{j=1}^{l_i} \mathbb{E} X_{ij} + \prod_{i=1}^k \mathbb{E} \left[\overline{X}_i^{s_i} \prod_{j=1}^{l_i} X_{ij} \right] + \frac{e}{2} \cdot \prod_{i=1}^K \mathbb{E} \left[\overline{X}_i^{2s_i} \right] \right).$$

We note that since $l_1 \vee l_2 \cdots \vee l_K \geq 1$,

$$\prod_{i=1}^K \prod_{j=1}^{l_i} \mathbb{E} X_{ij} = 0.$$

If $s_i + l_i$ is even for each $i \in [K]$, then, Lemma 15 gives us,

$$\prod_{i=1}^k \mathbb{E} \left[\overline{X}_i^{s_i} \prod_{j=1}^{l_i} X_{ij} \right] \leq \frac{C_k}{\sqrt{d^{k+\sum_{i=1}^K l_i}}}, \quad \prod_{i=1}^K \mathbb{E} \left[\overline{X}_i^{2s_i} \right] \leq \frac{C_k}{d^k}, \quad p_\pi(l_1, \dots, l_K) \leq \frac{C_k}{\sqrt{d^{k+\sum_{i=1}^K l_i}}}.$$

On the other hand, if $s_i + l_i$ is odd even for a single i , Lemma 15 gives us,

$$\prod_{i=1}^k \mathbb{E} \left[\overline{X}_i^{s_i} \prod_{j=1}^{l_i} X_{ij} \right] = 0, \quad \prod_{i=1}^K \mathbb{E} \left[\overline{X}_i^{2s_i} \right] \leq \frac{C_k}{d^k}, \quad p_\pi(l_1, \dots, l_K) \leq \frac{C_k}{d^k}.$$

Hence, when $l_i \leq s_i$ $i = 1, 2, \dots, K$, $l_1 \vee l_2 \cdots \vee l_K \geq 1$, we have shown,

$$p_\pi(l_1, \dots, l_K) \leq \begin{cases} C_k d^{-\frac{k+\sum_{i=1}^K l_i}{2}} & : l_i + s_i \text{ is even } \forall i \in [K] \\ C_k d^{-k} & : \text{otherwise} \end{cases}.$$

This concludes the analysis for this case.

Case 3: $\exists i : l_i > s_i$. By the case definition, we have,

$$\max_{i \in [K]} \left\lceil \frac{l_i}{s_i} \right\rceil \geq 2.$$

Without loss of generality, we can assume,

$$\max_{i \in [K]} \left\lceil \frac{l_i}{s_i} \right\rceil = \left\lceil \frac{l_1}{s_1} \right\rceil.$$

Let $t \geq 1$, $t \in \mathbb{N}$ be such that:

$$ts_1 < l_1 \leq (t+1)s_1.$$

By Lemma 8, we have,

$$p_\pi(l_1, l_2 \dots l_K) = \frac{1}{e} \cdot \mathbb{E} \left[e^{\bar{X}_1^{s_1} \dots \bar{X}_K^{s_K}} \prod_{i=1}^K \prod_{j=1}^{l_i} X_{ij} \right].$$

We use the bound:

$$1 + \sum_{a=1}^t \frac{x^a}{a!} - \frac{e|x|^{t+1}}{(t+1)!} \leq e^x \leq 1 + \sum_{a=1}^t \frac{x^a}{a!} + \frac{e|x|^{t+1}}{(t+1)!}, \quad |x| \leq 1.$$

Hence,

$$p_\pi(l_{1:K}) \leq \prod_{i=1}^K \prod_{j=1}^{l_i} \mathbb{E} X_{ij} + \sum_{a=1}^t \frac{1}{a!} \prod_{i=1}^K \mathbb{E} \left[\bar{X}_i^{as_i} \prod_{j=1}^{l_i} X_{ij} \right] + \frac{e}{(t+1)!} \prod_{i=1}^K \mathbb{E} [|\bar{X}_i|^{s_i(t+1)}].$$

Since $l_1 > s_1 \geq 0 \implies l_1 \geq 1$,

$$\prod_{i=1}^K \prod_{j=1}^{l_i} \mathbb{E} X_{ij} = 0.$$

For any $a \leq t$, $as_1 \leq ts_1 < l_1$. Hence by Lemma 15,

$$\mathbb{E} \left[\bar{X}_1^{as_1} \prod_{j=1}^{l_1} X_{1j} \right] = 0 \implies \sum_{a=1}^t \frac{1}{a!} \prod_{i=1}^K \mathbb{E} \left[\bar{X}_i^{as_i} \prod_{j=1}^{l_i} X_{ij} \right] = 0.$$

Finally again by Lemma 15,

$$p_\pi(l_1, l_2 \dots l_K) \leq \frac{e}{(t+1)!} \prod_{i=1}^K \mathbb{E} [|\bar{X}_i|^{s_i(t+1)}] \leq \frac{C_k}{d^{(t+1)k/2}}.$$

Note that,

$$(t+1) = \max_{i \in K} \left\lceil \frac{l_i}{s_i} \right\rceil \geq \max_{i \in K} \frac{l_i}{k} = \frac{l_1 \vee l_2 \dots \vee l_K}{k}.$$

Hence,

$$p_\pi(l_1, l_2 \dots l_K) \leq \frac{C_k}{d^{\frac{l_1 \vee l_2 \dots \vee l_K}{2}}}.$$

If $l_1 \vee l_2 \dots \vee l_K < 2k$, we can rely on the better bound (since $t \geq 1$):

$$p_\pi(l_1, l_2 \dots l_K) \leq \frac{C_k}{d^k}.$$

Hence, if $\exists i : l_i > s_i$, then,

$$p_\pi(l_1, l_2 \dots l_K) \leq \begin{cases} C_k \cdot d^{-k} & : l_1 \vee l_2 \dots \vee l_K < 2k \\ C_k \cdot d^{-\frac{l_1 \vee l_2 \dots \vee l_K}{2}} & : l_1 \vee l_2 \dots \vee l_K \geq 2k \end{cases}.$$

Combining all of these cases gives us the claim of the proposition. ■

Appendix D. Proof of Theorem 4

In this section we show how the various auxiliary results we have proved so far imply Theorem 4.

Proof In order to apply Proposition 13 we need to lower bound $\text{SDN}\left(D_0, \mathcal{D}(\pi), (3n)^{-\frac{1}{2}}\right)$. Lemma 14 gives us:

$$\begin{aligned} \text{SDN}\left(D_0, \mathcal{H}(\pi), \frac{1}{\sqrt{3n}}\right) &\geq \frac{1}{4\sqrt{3n}} \cdot \left(\sup_{q: \mathbb{E}_{D_0} q^2(\mathbf{T})=1} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \frac{1}{6\sqrt{n}} \right] \right)^{-1}. \end{aligned} \quad (18)$$

Let $u \geq 2$ be a integer parameter that we will set suitably later. Proposition 6 gives us, for any q such that $\mathbb{E}_{D_0} q^2(\mathbf{T}) = 1$,

$$\begin{aligned} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \frac{1}{6\sqrt{n}} \right] &\leq \left(36en \cdot \max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left((u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \right) \right)^{\frac{u}{2}}. \end{aligned}$$

In order to upper bound $\max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left((u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \right)$, we will utilize Proposition 9:

When $l_1 \vee l_2 \cdots \vee l_K \geq 2k$: Proposition C.2 gives us the bound:

$$(u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \leq \frac{C_k \cdot (u-1)^{kl_1 \vee l_2 \cdots \vee l_K}}{d^{l_1 \vee \dots \vee l_K / 2}} = \left(\frac{(u-1)^k}{\sqrt{d}} \right)^{l_1 \vee \dots \vee l_K}.$$

If $d \geq (u-1)^{2k}$, then the upper bound is a decreasing function of $l_1 \vee \dots \vee l_K$ and is maximized at $l_1 \vee \dots \vee l_K = 2k$. Hence,

$$\max_{l_1, \dots, l_K: l_1 \vee \dots \vee l_K \geq 2k} (u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \leq \frac{C_k (u-1)^{2k^2}}{d^k}, \quad \forall d \geq (u-1)^{2k}.$$

When $l_1 \vee l_2 \cdots \vee l_K < 2k$: Then,

$$\max_{l_1:K: l_1 \vee \dots \vee l_K \leq 2k} (u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \leq (u-1)^{2k^2} \cdot \max_{l_1:K: l_1 \vee \dots \vee l_K \leq 2k} p_\pi(l_{1:K}).$$

By Proposition 9, when $l_1 \vee \dots \vee l_K \leq 2k$, the upper bound on $p_\pi(l_{1:K})$ is maximized by setting:

$$l_i = \begin{cases} 0 & : s_i \text{ is even} \\ 1 & : s_i \text{ is odd} \end{cases}.$$

Recall that the oddness parameter \circ is precisely the number s_i that are odd, hence,

$$\max_{l_1, \dots, l_K: l_1 \vee \dots \vee l_K \leq 2k} (u-1)^{l_1+l_2+\dots+l_K} \cdot p_\pi(l_{1:K}) \leq \frac{C_k (u-1)^{2k^2}}{\sqrt{d^{k+\circ}}}.$$

Hence, we have obtained,

$$\max_{l_1, l_2, \dots, l_K \in \mathbb{N}_0} \left((u-1)^{l_1 + l_2 + \dots + l_K} \cdot p_\pi(l_1, l_2, \dots, l_K) \right) \leq \frac{C_k \cdot (u-1)^{2k^2}}{\sqrt{d^{k+\circ}}},$$

provided $d \geq (u-1)^{2k}$. Hence,

$$\sup_{q: \mathbb{E}_{D_0} q^2(\mathbf{T})=1} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \frac{1}{6\sqrt{n}} \right] \leq \left(\frac{C_k \cdot (u-1)^{2k^2} \cdot n}{\sqrt{d^{k+\circ}}} \right)^{\frac{u}{2}}$$

Substituting this in Eq. 18, we obtain,

$$\text{SDN} \left(D_0, \mathcal{H}(\pi), \frac{1}{\sqrt{3n}} \right) \geq \frac{1}{8\sqrt{n}} \cdot \left(\frac{1}{C_k \cdot (u-1)^{2k^2}} \cdot \frac{\sqrt{d^{k+\circ}}}{n} \right)^u, \quad (19)$$

provided $d \geq (u-1)^{2k}$. Now suppose there exist constants $C_0 \in (0, \infty)$ and $\epsilon \in (0, k)$ such that,

$$n \leq C_0 d^{\frac{k+\circ}{2} - \epsilon}.$$

Consider any $L \in \mathbb{N}$ and suppose that,

$$d \geq 11 \cdot \left(\frac{4L + k + \circ}{4\epsilon} \right)^{2k} \stackrel{\text{def}}{=} C_1(k, L, \epsilon). \quad (20)$$

We take $u = \epsilon^{-1} \cdot (L + (k + \circ)/4)$ in Eq. 19. This gives us,

$$\text{SDN} \left(D_0, \mathcal{H}(\pi), \frac{1}{\sqrt{3n}} \right) \geq c_1(k, L, \epsilon, C_0) \cdot d^L,$$

where,

$$c_1(k, L, \epsilon, C_0) \stackrel{\text{def}}{=} \frac{1}{8} \left(\frac{\epsilon}{(1 + C_0^2) \cdot C_k \cdot L} \right)^{\frac{2Lk^3}{\epsilon}}. \quad (21)$$

Hence, by Proposition 13, any SQ algorithm which solves the Tensor PCA testing problem with n samples must make at least $c_1(k, L, \epsilon C_0) \cdot d^L$ queries. In order to obtain an estimation lower bound, we appeal to Proposition 13 which lets us conclude that any SQ algorithm which solves the Tensor PCA estimation problem with n samples must make at least $0.5 \cdot c_1(k, L, \epsilon C_0) \cdot d^L$ queries.

In the case when $\circ = 0$, we can obtain a stronger estimation lower bound, by applying Proposition 13 with the base measure $\bar{D}(\pi)$. Lemma 14 and Proposition 6 tell us that $\text{SDN}(D_\star, \mathcal{H}(\pi), \epsilon)$ can be controlled in terms of the coefficients $\bar{p}_\pi(l_1, l_2, \dots, l_K)$. Furthermore from Proposition 9 that in general,

$$\bar{p}_\pi(l_1, \dots, l_K) \leq p_\pi(l_1, \dots, l_K),$$

but when $l_1 = l_2 \cdots = l_K = 0$, we have the improved upper bound $\bar{p}_\pi(0, 0 \dots, 0) \leq C_k \cdot d^{-k}$ (compared to $p_\pi(0, 0, \dots, 0) \leq C_k d^{-k/2}$). Consequently the quantity:

$$\max_{l_1, l_2 \dots l_K \in \mathbb{N}_0} \left((u-1)^{l_1+l_2 \dots l_K} \cdot \bar{p}_\pi(l_1, l_2 \dots l_K) \right),$$

is maximized by setting one $l_i = 2$ and the remaining ones as 0. This gives us,

$$\max_{l_1, l_2 \dots l_K \in \mathbb{N}_0} \left((u-1)^{l_1+l_2 \dots l_K} \cdot \bar{p}_\pi(l_1, l_2 \dots l_K) \right) \leq \frac{C_k \cdot (u-1)^{2k^2}}{\sqrt{d^{k+2}}},$$

which results in an improved lower bound:

$$\text{SDN}(D_\star, \mathcal{H}(\pi), \epsilon) \geq \frac{1}{8\sqrt{n}} \cdot \left(\frac{1}{C_k \cdot (u-1)^{2k^2}} \cdot \frac{\sqrt{d^{k+2}}}{n} \right)^u,$$

from which we can conclude via Proposition 13 that if $n \ll \sqrt{d^{k+2}}$, estimation is impossible in the SQ model with polynomially many queries. \blacksquare

Appendix E. Optimal Statistical Query Procedures

In this section we describe SQ procedures that achieve the optimal sample complexity for the Tensor PCA testing and estimation problems. These procedures leverage the “partial trace” technique. This approach was first introduced by Hopkins et al. (2016) who used this technique to compress matrices obtained by flattening higher order tensors in the context of symmetric Tensor PCA with $k = 3$ and some other planted inference problems. A different variant of this technique was used by Anandkumar et al. (2017) as a initialization procedure for tensor power method in the context of symmetric Tensor PCA with $k = 3$. Recent work by Biroli et al. (2019) generalized this approach for Symmetric Tensor PCA with arbitrary k .

We show that procedures based on the partial trace technique can be extended to Tensor PCA with partial symmetries and can be implemented in the SQ model via Fact 1. When implemented in the SQ model, these procedures suffer from a degradation in performance (when compared to the non-SQ performance on symmetric Tensor PCA problems from prior work (Hopkins et al., 2016; Anandkumar et al., 2017; Biroli et al., 2019)). However, the lower bounds of the previous section show that this is unavoidable.

We first design an SQ Test for even-order Symmetric Tensor PCA.

Lemma 17 (SQ Test for Even Symmetric Tensor PCA) *If $n \gg d^{\frac{k}{2}}$, there is a SQ algorithm which solves the Symmetric Tensor PCA testing problem with even k with $O(\log(n))$ queries.*

Proof Let $k = 2l$. Consider the query function:

$$q(\mathbf{T}) = \sum_{i_1, i_2 \dots i_l} T_{i_1, i_1, i_2, i_2 \dots i_l, i_l}.$$

It is easy to check that for any \mathbf{v} with $\|\mathbf{v}\| = \sqrt{d}$, under $D_{\mathbf{v}^{\otimes k}}$, $q(\mathbf{T}) \sim \mathcal{N}\left(1, d^{\frac{k}{2}}\right)$ where as, under D_0 , $q(\mathbf{T}) \sim \mathcal{N}\left(0, d^{\frac{k}{2}}\right)$. Hence for any $D \in \{D_0\} \cup \{D_{\mathbf{v}^{\otimes k}} : \mathbf{v} \in \mathbb{R}^d, \|\mathbf{v}\| = \sqrt{d}\}$,

$$\text{Var}_D q(\mathbf{T}) = 1, \mathbb{E}_D q^2(\mathbf{T}) \leq d^{\frac{k}{2}} + 1.$$

Consequently, we can use Fact 1 to construct an SQ procedure which returns an estimate of $\mathbb{E}_D q(\mathbf{T})$ denoted by $\hat{\mathbb{E}}_D q(\mathbf{T})$ such that,

$$|\hat{\mathbb{E}}_D q(\mathbf{T}) - \mathbb{E}_D q(\mathbf{T})| \rightarrow 0, \text{ as } d \rightarrow \infty,$$

if $n \gg d^{\frac{k}{2}}$. This requires only $O(\log(d))$ queries. This can be turned into a test for Tensor PCA by declaring $D = D_0$ if $\hat{\mathbb{E}}_D q(\mathbf{T}) \leq 0.5$ and $D \in \{D_{\mathbf{v}^{\otimes k}} : \mathbf{v} \in \mathbb{R}^d, \|\mathbf{v}\| = \sqrt{d}\}$ if $\hat{\mathbb{E}}_D q(\mathbf{T}) > 0.5$. The correctness of this test follows by the error guarantee of $\hat{\mathbb{E}}_D q(\mathbf{T})$. ■

Next we consider testing and estimation in generalized Tensor PCA with partial symmetries. We introduce the notion of the standard form of a labelling function.

Definition 18 *Recall that for a labelling function $\pi : [K] \rightarrow [k]$, we defined the oddness parameter of π (denoted by \circ) as the number of labels in $[K]$ which are used an odd number of times. Given a labelling function π , we can construct $l \stackrel{\text{def}}{=} (k - \circ)/2$ pairs from $k - \circ$ elements of $[k]$ by pairing two indices $i, j \in [k]$ if $\pi(i) = \pi(j)$. The remaining \circ indices are left unpaired. A labelling function π is in standard form if $\{2l + 1, 2l + 2, \dots, k\}$ are the unpaired indices and,*

$$\pi(2i - 1) = \pi(2i), i \in [l].$$

If π is in standard form, then for any $D \in \mathcal{D}(\pi)$,

$$\mathbb{E}_D \mathbf{T} = \mathbf{E}_1 \otimes \mathbf{E}_2 \cdots \otimes \mathbf{E}_l \otimes \mathbf{O}$$

where $\mathbf{E}_1, \dots, \mathbf{E}_l$ are rank-1 symmetric matrices with $\|\mathbf{E}_i\| = \text{Tr}(\mathbf{E}_i) = 1$ and \mathbf{O} is a \circ -order rank 1 tensor with $\|\mathbf{O}\| = 1$.

Lemma 19 *There exists a SQ algorithm which given the knowledge of the labelling function π solves the Tensor PCA estimation problem after making at most $O(d^k \cdot \log(n))$ queries provided the sample size n satisfies:*

$$n \gg \begin{cases} \sqrt{d^{k+2}} & : \circ = 0 \\ \sqrt{d^{k+\circ}} & : \circ \geq 1 \end{cases}.$$

Furthermore when $\circ \geq 1$, there is a test based on this estimator that solves Tensor PCA testing problem under the same assumptions on the sample size n .

Proof Since the algorithm knows the labelling function π , we can assume, without loss of generality that π is in standard form since by applying a suitable permutation to the modes of the tensor, the algorithm can always convert the Tensor PCA problem to the standard

form. We define $l \stackrel{\text{def}}{=} (k - \mathfrak{o})/2$, where \mathfrak{o} is the number of labels used an odd number of times by π . We first consider the Tensor PCA estimation problem. Recall that by the definition of standard form, for any $D \in \mathcal{D}(\pi)$,

$$\mathbb{E}_D \mathbf{T} = \mathbf{E}_1 \otimes \mathbf{E}_2 \cdots \otimes \mathbf{E}_l \otimes \mathbf{O}, \quad \forall D \in \mathcal{D}(\pi).$$

Estimating \mathbf{O} : We construct an estimator for \mathbf{O} by estimating it entry-wise. Consider the query functions:

$$q_{i_1, i_2 \dots i_{\mathfrak{o}}}(\mathbf{T}) = \sum_{j_1, j_2 \dots, j_l} \mathbf{T}_{j_1, j_1, j_2, j_2 \dots j_l, j_l, i_1, i_2 \dots i_{\mathfrak{o}}}.$$

For any $D \in \mathcal{D}(\pi)$

$$q_{i_1, i_2 \dots i_{\mathfrak{o}}}(\mathbf{T}) \sim \mathcal{N}\left(\text{Tr}(\mathbf{E}_1) \cdot \text{Tr}(\mathbf{E}_2) \cdots \text{Tr}(\mathbf{E}_l) O_{i_1, i_2 \dots i_{\mathfrak{o}}}, d^l\right) = \mathcal{N}\left(O_{i_1, i_2 \dots i_{\mathfrak{o}}}, d^l\right).$$

Hence using Fact 1, one can obtain estimates such that,

$$\left| \hat{\mathbb{E}}_D q_{i_1, i_2 \dots i_{\mathfrak{o}}}(\mathbf{T}) - \mathbb{E}_D q_{i_1, i_2 \dots i_{\mathfrak{o}}}(\mathbf{T}) \right| \leq 9 \cdot \log(n) \cdot \sqrt{\frac{d^l}{n}}, \quad \forall D \in \mathcal{D}(\pi).$$

Now we construct an estimate of \mathbf{O} , denoted by $\hat{\mathbf{O}}$ with $\hat{O}_{i_1, i_2 \dots i_{\mathfrak{o}}} = \hat{\mathbb{E}}_D q_{i_1, i_2 \dots i_{\mathfrak{o}}}(\mathbf{T})$. This estimate has the error guarantee:

$$\begin{aligned} \|\hat{\mathbf{O}} - \mathbf{O}\|^2 &= \sum_{i_1, i_2 \dots i_{\mathfrak{o}}} (\hat{O}_{i_1, i_2 \dots i_{\mathfrak{o}}} - O_{i_1, i_2 \dots i_{\mathfrak{o}}})^2 \\ &\leq 81 \cdot \log^2(n) \cdot \frac{d^{l+\mathfrak{o}}}{n} = 81 \cdot \log^2(n) \cdot \frac{\sqrt{d^{k+\mathfrak{o}}}}{n}. \end{aligned}$$

Under the sample size assumption, $\|\hat{\mathbf{O}} - \mathbf{O}\| \rightarrow 0$ as $d \rightarrow \infty$. This further implies $\|\hat{\mathbf{O}}\| \rightarrow \|\mathbf{O}\| = 1$ and, $\langle \hat{\mathbf{O}}, \mathbf{O} \rangle \rightarrow 1$. Note that when $\mathfrak{o} = 0$, we can simply skip this step.

Estimating $\mathbf{E}_{1:l}$ We estimate \mathbf{E}_1 entry-wise by using the queries:

$$q_{ab}(\mathbf{T}) = \frac{1}{\|\hat{\mathbf{O}}\|} \cdot \sum_{j_1, j_2, \dots, j_{l-2}, i_1, i_2 \dots i_{\mathfrak{o}}} T_{a, b, j_1, j_1, j_2, j_2 \dots, j_{l-1}, j_{l-1}, i_1, i_2 \dots i_{\mathfrak{o}}} \cdot \hat{O}_{i_1, i_2 \dots i_{\mathfrak{o}}}, \quad a, b \in [d].$$

Under $D \in \mathcal{D}(\pi)$,

$$q_{ab}(\mathbf{T}) \sim \mathcal{N}\left((\mathbf{E}_1)_{a,b} \cdot \frac{\langle \mathbf{O}, \hat{\mathbf{O}} \rangle}{\|\hat{\mathbf{O}}\|}, d^{l-1}\right).$$

Hence using Fact 1, we can obtain estimates $\hat{\mathbb{E}}_D q_{ab}(\mathbf{T})$ with the guarantee:

$$\left| \hat{\mathbb{E}}_D q_{ab}(\mathbf{T}) - \mathbb{E}_D q_{ab}(\mathbf{T}) \right| \leq 9 \cdot \log(n) \cdot \sqrt{\frac{d^{l-1}}{n}}, \quad \forall D \in \mathcal{D}(\pi).$$

We construct an estimator of \mathbf{E}_1 , denoted by $\hat{\mathbf{E}}_1$ with entries defined by $(\hat{\mathbf{E}}_1)_{a,b} = \hat{\mathbb{E}}_D q_{ab}(\mathbf{T})$. The error guarantee of this estimator is given by:

$$\left\| \hat{\mathbf{E}}_1 - \frac{\langle \mathbf{O}, \hat{\mathbf{O}} \rangle}{\|\hat{\mathbf{O}}\|} \mathbf{E}_1 \right\|^2 = 81 \cdot \log^2(n) \cdot \frac{d^{l+1}}{n} = 81 \cdot \log^2(n) \cdot \frac{\sqrt{d^{k-\circ+2}}}{n}.$$

By Triangle Inequality,

$$\|\hat{\mathbf{E}}_1 - \mathbf{E}_1\| \leq \left\| \hat{\mathbf{E}}_1 - \frac{\langle \mathbf{O}, \hat{\mathbf{O}} \rangle}{\|\hat{\mathbf{O}}\|} \mathbf{E}_1 \right\| + \left| 1 - \frac{\langle \mathbf{O}, \hat{\mathbf{O}} \rangle}{\|\hat{\mathbf{O}}\|} \right| = 9 \cdot \log(n) \cdot \left(\frac{\sqrt{d^{k-\circ+2}}}{n} \right)^{\frac{1}{2}} + o_d(1).$$

When $\circ = 0$, the sample size assumption $n \gg \sqrt{d^{k+2}}$ guarantees $\|\mathbf{E}_1 - \hat{\mathbf{E}}_1\| \rightarrow 0$ as $d \rightarrow \infty$. When $\circ \geq 1$ since $k - \circ + 2 \leq k + \circ$, the assumption $n \gg \sqrt{d^{k+\circ}} \implies n \gg \sqrt{d^{k-\circ+2}}$ and hence, $\|\mathbf{E}_1 - \hat{\mathbf{E}}_1\| \rightarrow 0$. By a similar procedure, we can also estimate $\mathbf{E}_2, \mathbf{E}_3 \dots \mathbf{E}_l$. Hence we have obtained estimators $\hat{\mathbf{E}}_{1:l}$ such that,

$$\|\hat{\mathbf{E}}_i - \mathbf{E}_i\| \rightarrow 0, \|\hat{\mathbf{E}}_i\| \rightarrow 1, \langle \mathbf{E}_i, \hat{\mathbf{E}}_i \rangle \rightarrow 1 \text{ as } d \rightarrow \infty.$$

Estimating $\mathbb{E}_D q(\mathbf{T})$ We construct an estimator of $\mathbb{E}_D \mathbf{T}$ as:

$$\hat{\mathbb{E}}_D \mathbf{T} = \frac{\hat{\mathbf{E}}_1}{\|\hat{\mathbf{E}}_1\|} \otimes \frac{\hat{\mathbf{E}}_2}{\|\hat{\mathbf{E}}_2\|} \dots \otimes \frac{\hat{\mathbf{E}}_l}{\|\hat{\mathbf{E}}_l\|} \otimes \frac{\hat{\mathbf{O}}}{\|\hat{\mathbf{O}}\|}.$$

The error of this estimator can be computed as:

$$\|\hat{\mathbb{E}}_D \mathbf{T} - \mathbb{E}_D \mathbf{T}\|^2 = 1 + 1 - 2 \cdot \left(\prod_{i=1}^l \frac{\langle \hat{\mathbf{E}}_i, \mathbf{E}_i \rangle}{\|\hat{\mathbf{E}}_i\|} \right) \cdot \frac{\langle \hat{\mathbf{O}}, \mathbf{O} \rangle}{\|\hat{\mathbf{O}}\|} \rightarrow 0 \text{ as } d \rightarrow \infty.$$

This proves the claim of the lemma for estimation. For testing we observe that we only need to consider the case $\circ \geq 1$ since Lemma 17 already handles the case $\circ = 0$. Hence for any $D \in \mathcal{D}(\pi)$, the estimator $\hat{\mathbf{O}}$ satisfies $\|\hat{\mathbf{O}}\| \rightarrow 1$. When $D = D_0$, the same analysis can be repeated to show $\|\hat{\mathbf{O}}\| \rightarrow 0$. Hence we construct a test by declaring $D = D_0$ if $\|\hat{\mathbf{O}}\| \leq 0.5$ and $D \in \mathcal{D}(\pi)$ if $\|\hat{\mathbf{O}}\| > 0.5$. This test solves the Tensor PCA testing problem for large enough d . \blacksquare

Appendix F. Poisson Random Tensors

This section is devoted to the proof of Proposition 7 which describes a certain conditional independence structure in Poisson random tensors. We first collect some standard facts regarding the Multinomial and the Poisson distributions which will be useful for our analysis.

Definition 20 (Multinomial Distribution) For any $n \in \mathbb{N}, k \in \mathbb{N}$, the Multinomial Distribution with parameters (n, k) (denoted by $\text{Mult}(n, k)$) is the distribution supported on the set:

$$\text{Supp}(\text{Mult}(n, k)) = \left\{ \mathbf{x} \in \mathbb{R}^k : \forall i \in [k], x_i \in \mathbb{N}_0, \sum_{i=1}^k x_i = n \right\},$$

with the probability mass function given by:

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_k = x_k) = \binom{n}{x_1, x_2, \dots, x_k} \cdot \frac{1}{k^n}, \forall \mathbf{x} \in \text{Supp}(\text{Mult}(n, k)).$$

Fact 2 (Properties of the Multinomial Distribution) Let $\mathbf{X} \sim \text{Mult}(n, k)$ and $\mathbf{Y} \sim \text{Mult}(m, k)$ be two independent Multinomial random vectors. Then we have,

1. The generating function of \mathbf{X} is given by:

$$\mathbb{E} \left[z_1^{X_1} z_2^{X_2} \dots z_k^{X_k} \right] = \left(\frac{\sum_{i=1}^k z_i}{k} \right)^n, \forall \mathbf{z} \in \mathbb{C}^k.$$

2. The sum $\mathbf{X} + \mathbf{Y}$ is also a Multinomial random vector with $\mathbf{X} + \mathbf{Y} \sim \text{Mult}(n + m, k)$.

Definition 21 (Poisson Distribution) The Poisson distribution with parameter $\lambda \geq 0$ denoted by $\text{Pois}(\lambda)$ is the distribution on \mathbb{N}_0 with the probability mass function:

$$\mathbb{P}(X = x) = \frac{e^{-\lambda} \lambda^x}{x!}.$$

Fact 3 (Properties of the Poisson Distribution) Let X, X_1, X_2, \dots, X_k be i.i.d. $\text{Pois}(\lambda)$ distributed random variables. Then,

1. Generating Function of $\text{Pois}(\lambda)$: We have,

$$\mathbb{E} z^X = e^{\lambda(z-1)}, \forall z \in \mathbb{C}.$$

2. The sum $T = \sum_{i=1}^k X_i$ is also a Poisson random variable with $T \sim \text{Pois}(k\lambda)$.
3. The conditional distribution of the random vector (X_1, X_2, \dots, X_k) given T is multinomial:

$$(X_1, X_2, \dots, X_k) \mid T \sim \text{Mult}(T, k).$$

We now prove the conditional independence property of Poisson tensors (Proposition 7) which we have restated below for convenience. Let \mathbf{C} be a random tensor in $\otimes^k \mathbb{W}^d$ whose entries are sampled independently as follows:

$$C_{i_1, i_2, \dots, i_k} \stackrel{\text{i.i.d.}}{\sim} \text{Pois} \left(\frac{1}{d^k} \right).$$

Proposition 22 *Let \mathbf{C} be a Poisson random tensor sampled as above. Then, we have*

$$\|\mathbf{C}\|_1 \sim \text{Pois}(1).$$

Moreover, conditioned on $\|\mathbf{C}\|_1$, we have

$$\text{SUM}(\mathbf{C}, i) \mid \|\mathbf{C}\|_1 \stackrel{i.i.d.}{\sim} \text{Mult}(\|\mathbf{C}\|_1, d), \forall i \in [k].$$

Proof The claim that $\|\mathbf{C}\|_1 \sim \text{Pois}(1)$ is a direct consequence of the additivity property of the Poisson Distribution (cf. Fact 3.3). We will prove the claim about the conditional distribution of $\text{SUM}(\mathbf{C}, i)$ given T by induction on the order of the tensor, k .

Base Case, $k = 2$: Note that in this case \mathbf{C} is a $d \times d$ matrix. $\text{SUM}(\mathbf{C}, 1)$ is the vector of row-wise sums and $\text{SUM}(\mathbf{C}, 2)$ are the column wise sums. We first characterise the distribution of $\text{SUM}(\mathbf{C}, 1) \mid \|\mathbf{C}\|_1$. Due to the additivity property of the Poisson distribution (cf. Fact 3.2),

$$\text{SUM}(\mathbf{C}, 1)_i \stackrel{i.i.d.}{\sim} \text{Pois}\left(\frac{1}{d}\right), \forall i \in [d].$$

Since $\|\mathbf{C}\|_1 = \sum_{i=1}^d \text{SUM}(\mathbf{C}, 1)_i$, appealing to the Poisson-Multinomial connection (cf. Fact 3.3), we have,

$$\text{SUM}(\mathbf{C}, 1) \mid \|\mathbf{C}\|_1 \sim \text{Mult}(\|\mathbf{C}\|_1, d).$$

Next we characterize the conditional distribution of $\text{SUM}(\mathbf{C}, 2) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1)$. We first observe that,

$$\begin{aligned} (C_{i,1}, C_{i,2} \dots C_{i,d}) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1) &\stackrel{d}{=} (C_{i,1}, C_{i,2} \dots C_{i,n}) \mid \text{SUM}(\mathbf{C}, 1)_i \\ &\sim \text{Mult}(\text{SUM}(\mathbf{C}, 1)_i, d). \end{aligned}$$

In the last step, we again appealed to the Poisson-Multinomial connection (Fact 3.3). Hence by the additivity property of the Multinomial distribution (Fact 2.2), we obtain,

$$\begin{aligned} \text{SUM}(\mathbf{C}, 2) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1) &= \sum_{i=1}^n (C_{i,1}, C_{i,2} \dots C_{i,n}) \Bigg| \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1) \\ &\sim \text{Mult}\left(\sum_{i=1}^n \text{SUM}(\mathbf{C}, 1)_i, d\right) \\ &= \text{Mult}(\|\mathbf{C}\|_1, d). \end{aligned}$$

Hence,

$$\text{SUM}(\mathbf{C}, 1) \mid \|\mathbf{C}\|_1 \sim \text{Mult}(\|\mathbf{C}\|_1, d), \text{SUM}(\mathbf{C}, 2) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1) \sim \text{Mult}(\|\mathbf{C}\|_1, d),$$

which proves the proposition for $k = 2$.

Induction Hypothesis: We assume the claim of the proposition holds for Poisson tensors of order $k - 1$.

Induction Step: We now prove the proposition for Tensors of order k . Since we plan to appeal to the induction hypothesis, we define a $k - 1$ order tensor $\tilde{\mathbf{C}}$ by summing over mode k of \mathbf{C} :

$$\tilde{C}_{j_1, j_2, \dots, j_{k-1}} = \sum_{j_k} C_{j_1, j_2, \dots, j_k}, \quad \forall j_1, j_2, \dots, j_{k-1} \in [d].$$

We note that,

$$\|\mathbf{C}\|_1 = \|\tilde{\mathbf{C}}\|_1, \quad \text{SUM}(\mathbf{C}, i) = \text{SUM}(\tilde{\mathbf{C}}, i), \quad \forall i \in [k - 1].$$

Hence by the induction hypothesis, we have,

$$\text{SUM}(\mathbf{C}, i) \mid \|\mathbf{C}\|_1 \stackrel{\text{i.i.d.}}{\sim} \text{Mult}(\|\mathbf{C}\|_1, d), \quad \forall i \in [k - 1].$$

We now just need to characterize the distribution of

$$\text{SUM}(\mathbf{C}, k) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1), \text{SUM}(\mathbf{C}, 2) \dots \text{SUM}(\mathbf{C}, k - 1).$$

We first study its distribution conditioned on a larger σ -algebra: $\text{SUM}(\mathbf{C}, k) \mid \tilde{\mathbf{C}}$. Observe that,

$$\begin{aligned} (C_{j_1, \dots, j_{k-1}, 1}, C_{j_1, \dots, j_{k-1}, 2}, \dots, C_{j_1, \dots, j_{k-1}, d}) \Big| \tilde{\mathbf{C}} \\ \stackrel{d}{=} (C_{j_1, j_2, \dots, j_{k-1}, 1}, C_{j_1, j_2, \dots, j_{k-1}, 2}, \dots, C_{j_1, j_2, \dots, j_{k-1}, d}) \Big| \tilde{C}_{j_1, \dots, j_{k-1}} \\ \sim \text{Mult}(\tilde{C}_{j_1, j_2, \dots, j_{k-1}}, d). \end{aligned}$$

In the last step we used the Multinomial-Poisson connection (Fact 3.3). Hence, by additivity of Multinomial distributions (Fact 2.2),

$$\begin{aligned} \text{SUM}(\mathbf{C}, k) \mid \tilde{\mathbf{C}} &= \sum_{j_1, j_2, \dots, j_{k-1}} (C_{j_1, \dots, j_{k-1}, 1}, C_{j_1, \dots, j_{k-1}, 2}, \dots, C_{j_1, \dots, j_{k-1}, d}) \Big| \tilde{\mathbf{C}} \\ &\sim \text{Mult}\left(\sum_{j_1, j_2, \dots, j_{k-1}} \tilde{C}_{j_1, j_2, \dots, j_{k-1}}, d\right) \\ &= \text{Mult}(\|\mathbf{C}\|_1, d). \end{aligned}$$

This shows that,

$$\text{SUM}(\mathbf{C}, k) \mid \|\mathbf{C}\|_1, \text{SUM}(\mathbf{C}, 1), \text{SUM}(\mathbf{C}, 2) \dots \text{SUM}(\mathbf{C}, k - 1) \sim \text{Mult}(\|\mathbf{C}\|_1, d).$$

This concludes the induction step and the proof of the proposition. ■

Appendix G. Proof of Proposition 13 and Lemma 14

This section is dedicated to missing proofs from Appendix A on the framework for proving SQ lower bounds.

G.1 Proof of Proposition 13

Proof As mentioned previously, the claim regarding the testing problem is simply Theorem 7.1 of Feldman et al. (2018). The proof of the claim for estimation is a modification of their proof. We define B as:

$$B \stackrel{\text{def}}{=} \text{SDN}\left(D_\star, \mathcal{H}(\pi), (3n)^{-\frac{1}{2}}\right).$$

We will prove the contrapositive of this statement: For any SQ algorithm \mathcal{A} which makes $t < B/2$ queries, we will exhibit a $D \in \mathcal{H}(\pi)$ and a corresponding $\text{VSTAT}_D(n)$ oracle which when used to answer queries made by \mathcal{A} results in \mathcal{A} outputting an estimator $\hat{\mathbf{V}}$ whose error satisfies $\|\hat{\mathbf{V}} - \mathbb{E}_D q(\mathbf{T})\| \geq 0.5$ and hence \mathcal{A} fails to solve the Tensor PCA estimation in the sense of Definition 3.

We begin by specifying the adversarial $\text{VSTAT}_D(n)$ oracle: This oracle avoids committing to a specific $D \in \mathcal{H}(\pi)$ in the beginning and simply responds to any query q by the estimate $\mathbb{E}_{D_\star} q(\mathbf{T})$. After the learning algorithm has exhausted its query budget by asking a sequence of queries q_1, q_2, \dots, q_t and declares an estimator $\hat{\mathbf{V}}$, the oracle chooses a $D \in \mathcal{H}(\pi)$ to maximize the error of the estimate subject to the constraint that the responses are valid $\text{VSTAT}_D(n)$ responses. Define the sets:

$$A_i \stackrel{\text{def}}{=} \left\{ (\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) : \mathbf{v}_j \in \{\pm 1\}^d, |p_i(\mathbf{v}_{1:K}) - \mathbb{E}_{D_\star} q_i(\mathbf{T})| > \frac{1}{n} \vee \sqrt{\frac{p_i(\mathbf{v}_{1:K}) - p_i^2(\mathbf{v}_{1:K})}{n}} \right\},$$

where,

$$p_i(\mathbf{v}_1, \mathbf{v}_2 \cdots \mathbf{v}_K) \stackrel{\text{def}}{=} \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}}} q_i(\mathbf{T}).$$

Then by the definition of a $\text{VSTAT}_D(n)$ oracle (Definition 1), for any D such that:

$$D \in \mathcal{H}(\pi) \setminus \bigcup_{i=1}^t A_i,$$

the above specified adversarial oracle is a valid $\text{VSTAT}_D(n)$ oracle. We make the following claims:

Claim 1: For all $i \in [t]$, $|A_i| \leq 2^{dK}/B$.

Claim 2: There exist D_1, D_2 such that, $D_1, D_2 \in \mathcal{H}(\pi) \setminus \bigcup_{i=1}^t A_i$ & $\|\mathbb{E}_{D_1} \mathbf{T} - \mathbb{E}_{D_2} \mathbf{T}\| \geq 1$.

We first assume the above claims are true and prove Proposition 13. By the triangle inequality,

$$\|\mathbb{E}_{D_1} \mathbf{T} - \mathbb{E}_{D_2} \mathbf{T}\| \leq \|\mathbb{E}_{D_1} \mathbf{T} - \hat{\mathbf{V}}\| + \|\mathbb{E}_{D_2} \mathbf{T} - \hat{\mathbf{V}}\| \leq 2 \max(\|\mathbb{E}_{D_1} \mathbf{T} - \hat{\mathbf{V}}\|, \|\mathbb{E}_{D_2} \mathbf{T} - \hat{\mathbf{V}}\|).$$

Appealing to Claim 2, we obtain,

$$\max(\|\mathbb{E}_{D_1} \mathbf{T} - \hat{\mathbf{V}}\|, \|\mathbb{E}_{D_2} \mathbf{T} - \hat{\mathbf{V}}\|) \geq \frac{1}{2}.$$

Hence, we have found a $D \in \mathcal{D}(\pi)$ and a valid $\text{VSTAT}_D(n)$ oracle such that,

$$\|\hat{\mathbf{V}} - \mathbb{E}_D \mathbf{T}\| \geq \frac{1}{2},$$

which was the claim of Proposition 2. Finally we provide a proof of the claims made above.

Proof of Claim 1: Exactly the same claim has been proved in the proof of Theorem 7.1 in Feldman et al. (2018), so we refer the reader to this paper for the proof.

Proof of Claim 2: As the SQ algorithm makes more and more queries, the set of valid choices of D for the oracle becomes smaller. We can track the set of valid choices of D and the relative distances between them by a sequence of undirected graphs. The initial graph $G_0 = (V_0, E_0)$ has the vertex set

$$V_0 = \{(\mathbf{v}_1, \mathbf{v}_2 \cdots, \mathbf{v}_K) : \mathbf{v}_j \in \{\pm 1\}^d, \forall j \in [K]\},$$

and the edge set:

$$E_0 = \{[(\mathbf{u}_1, \dots, \mathbf{u}_K), (\mathbf{v}_1, \dots, \mathbf{v}_K)] : \mathbf{u}_i, \mathbf{v}_i \in \{\pm 1\}^d, |\langle \mathbf{u}_i, \mathbf{v}_i \rangle| \leq 2^{-\frac{1}{k}} \cdot d \forall i \in [K]\}.$$

Note that for any two connected vertices $(\mathbf{u}_1, \dots, \mathbf{u}_K)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_K)$:

$$\begin{aligned} \|\mathbb{E}_{D_{\mathbf{u}_{\pi(1)} \cdots \mathbf{u}_{\pi(k)}}} \mathbf{T} - \mathbb{E}_{D_{\mathbf{v}_{\pi(1)} \cdots \mathbf{v}_{\pi(k)}}} \mathbf{T}\|^2 &= \frac{1}{d^k} \|\mathbf{u}_{\pi(1)} \cdots \otimes \mathbf{u}_{\pi(k)} - \mathbf{v}_{\pi(1)} \cdots \otimes \mathbf{v}_{\pi(k)}\|^2 \\ &= 2 - 2 \prod_{i=1}^k \frac{\langle \mathbf{u}_{\pi(i)}, \mathbf{v}_{\pi(i)} \rangle}{d} \\ &\geq 2 - 2 \cdot \left(\frac{1}{2^{\frac{1}{k}}}\right)^k = 1. \end{aligned}$$

Due to the underlying symmetry between the vertices, it is easy to check that G_0 is a r -regular graph for some $r \geq 1$. Hence $|E_0| = |V_0| \cdot r/2 = r \cdot 2^{Kd-1}$. After every query q_t , we update the graph as follows: The new graph $G_i = (V_i, E_i)$ where $V_i = V_{i-1} \setminus A_i$ and E_i is the set of remaining edges after all edges which involved a vertex in A_i are removed from E_{i-1} . Since each vertex has at most r edges, we have,

$$|E_{i-1}| - |E_i| \leq r|A_i| \stackrel{\text{Claim 1}}{\leq} \frac{r \cdot 2^{dK}}{B}.$$

Hence,

$$|E_0| - |E_t| \leq \frac{r \cdot t \cdot 2^{dK}}{B}$$

Since $t < B/2$, $|E_t| > |E_0| - r \cdot 2^{dK-1} = 0$. Hence $|E_t| \geq 1$, that is, the graph after t queries contains at least one remaining edge. The vertices corresponding to this edge yield distributions D_1, D_2 satisfying Claim 2. \blacksquare

G.2 Proof of Lemma 14

Proof In order to show the desired lower bound on $\text{SDN}(D_\star, \mathcal{H}(\pi), \epsilon)$, it is sufficient to verify the implication in Eq. 12 from Definition 12 with

$$m = \frac{\epsilon}{4} \cdot \left(\sup_{q: \mathbb{E}_{D_0} q^2(\mathbf{T})=1} \mathbb{P}_{D \sim \text{Unif}(\mathcal{H})} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| > \frac{\epsilon}{2} \right] \right)^{-1}.$$

To this goal consider any $\mathcal{S} \subset \mathcal{H}(\pi)$ with $|\mathcal{S}| \geq m^{-1} \cdot |\mathcal{H}(\pi)|$ and any q with $\mathbb{E}_{D_\star} q^2(\mathbf{T}) = 1$. First we observe that,

$$\begin{aligned} \frac{|D \in \mathcal{S} : |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| > \frac{\epsilon}{2}|}{|\mathcal{S}|} &\leq \frac{|D \in \mathcal{H}(\pi) : |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| > \frac{\epsilon}{2}|}{|\mathcal{S}|} \\ &= \frac{|\mathcal{H}(\pi)|}{|\mathcal{S}|} \cdot \mathbb{P}_{D \sim \text{Unif}(\mathcal{H}(\pi))} \left[|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| > \frac{\epsilon}{2} \right] \\ &\leq \frac{\epsilon}{4}. \end{aligned} \quad (22)$$

Consider any $D \in \mathcal{H}(\pi)$. Hence, there exist hypercube vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K \in \{\pm 1\}^d$ such that $D = D_{\mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)}}$. We first prove a worst case upper bound on $|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})|$ when $D_\star = D_0$:

$$\begin{aligned} |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| &= \left| \mathbb{E}_{D_0} \left(\frac{d\mathbb{P}_D}{d\mathbb{P}_{D_0}}(\mathbf{T}) - 1 \right) q(\mathbf{T}) \right| \\ &\stackrel{(a)}{\leq} \sqrt{\mathbb{E}_{D_0} q^2(\mathbf{T})} \cdot \sqrt{\mathbb{E}_{D_0} \left(\frac{d\mathbb{P}_D}{d\mathbb{P}_{D_0}}(\mathbf{T}) - 1 \right)^2} \\ &= \sqrt{\mathbb{E}_{D_0} \left(\frac{d\mathbb{P}_D}{d\mathbb{P}_{D_0}}(\mathbf{T}) \right)^2} - 1 \\ &= \left(\mathbb{E}_{D_0} \exp \left(\|\mathbf{T}\|^2 - \|\mathbf{T} - d^{-\frac{k}{2}} \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)}\|^2 \right) - 1 \right)^{\frac{1}{2}} \\ &= \left(\mathbb{E}_{D_0} \exp \left(2 \langle \mathbf{T}, d^{-\frac{k}{2}} \mathbf{v}_{\pi(1)} \otimes \mathbf{v}_{\pi(2)} \cdots \otimes \mathbf{v}_{\pi(k)} \rangle - 1 \right) - 1 \right)^{\frac{1}{2}} \\ &\stackrel{(b)}{=} \sqrt{e - 1} \leq 2. \end{aligned}$$

In the above display, in the equation marked (a), we applied Cauchy-Schwarz inequality and in the step marked (b), we noted that under D_0 , the entries of \mathbf{T} are i.i.d. $\mathcal{N}(0, 1)$ and used the formula for the Gaussian moment generating function. Hence we have obtained, for any $D \in \mathcal{H}(\pi)$

$$|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_0} q(\mathbf{T})| \leq 2. \quad (23)$$

An exactly analogous calculation shows that:

$$|\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{\bar{D}(\pi)} q(\mathbf{T})| \leq \sqrt{e - 1} \leq 2.$$

We can now verify the implication in Eq. 12:

$$\begin{aligned} \frac{\sum_{D \in \mathcal{S}} |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})|}{|\mathcal{S}|} &\leq \frac{\epsilon}{2} + \frac{\sum_{D : |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})| > \frac{\epsilon}{2}} |\mathbb{E}_D q(\mathbf{T}) - \mathbb{E}_{D_\star} q(\mathbf{T})|}{|\mathcal{S}|} \\ &\stackrel{(a)}{\leq} \frac{\epsilon}{2} + \frac{\epsilon}{4} \cdot 2 \\ &\leq \epsilon. \end{aligned}$$

In the above display, the step marked (a) is a consequence of Eqs. 22 and 23. This verifies the implication in Eq. 12 and completes the proof of this lemma. \blacksquare

Appendix H. Fourier Analytic Tools

This section collects some basic facts from Fourier analysis on the Boolean hypercube $\{\pm 1\}^d$ and the Gaussian Hilbert space $\mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d))$. Our reference for all these results is the book of O’Donnell (2014).

H.1 Fourier Analysis on $\{\pm 1\}^d$

Consider the class of arbitrary functions $f : \{\pm 1\}^d \mapsto \mathbb{R}$:

$$\mathcal{F}_{\text{BOOLEAN}}(d) = \{f : \{\pm 1\}^d \mapsto \mathbb{R}\}.$$

The following is a orthonormal basis for $\mathcal{F}_{\text{BOOLEAN}}(d)$:

$$\text{BASIS}(\mathcal{F}_{\text{BOOLEAN}}(d)) = \left\{ \mathbf{z}^{\mathbf{r}} \stackrel{\text{def}}{=} \prod_{i=1}^d z_i^{r_i} : \mathbf{r} \in \{0, 1\}^d \right\}.$$

The above basis is orthonormal in the following sense:

$$\mathbb{E}_{\text{Unif}(\pm 1)^d} \mathbf{Z}^{\mathbf{r}} \mathbf{Z}^{\mathbf{s}} = \begin{cases} 0 & \mathbf{r} \neq \mathbf{s} \\ 1 & \mathbf{r} = \mathbf{s} \end{cases}.$$

Hence given any function $f \in \mathcal{F}_{\text{BOOLEAN}}(d)$ has a unique representation of the form:

$$f(\mathbf{z}) = \sum_{\mathbf{r} \in \{0, 1\}^d} \hat{f}(\mathbf{r}) \mathbf{z}^{\mathbf{r}}.$$

The coefficients $\hat{f}(\mathbf{r})$ are called the Fourier coefficients of f . This representation satisfies the usual Parseval’s identity:

$$\mathbb{E}_{\text{Unif}(\{\pm 1\}^d)} f^2(\mathbf{Z}) = \sum_{\mathbf{r} \in \{0, 1\}^d} \hat{f}^2(\mathbf{r}).$$

For any $\lambda \geq 1$, we define the operator $\text{SMOOTH}_{\lambda}[\cdot] : \mathcal{F}_{\text{BOOLEAN}}(d) \mapsto \mathcal{F}_{\text{BOOLEAN}}(d)$ as follows:

$$\text{SMOOTH}_{\lambda}[f](\mathbf{z}) \stackrel{\text{def}}{=} \sum_{\mathbf{r} \in \{0, 1\}^d} \hat{f}(\mathbf{r}) \cdot \left(\frac{1}{\lambda}\right)^{\|\mathbf{r}\|_1} \cdot \mathbf{z}^{\mathbf{r}}.$$

The operator $\text{SMOOTH}_{\lambda}[\cdot]$ smooths a function by down weighting the high degree Fourier coefficients. Hypercontractivity Theorems control L_q norm of the random variable $\text{SMOOTH}_{\lambda}[f](\mathbf{Z})$ in terms of L_p norm of $f(\mathbf{Z})$ where $\mathbf{Z} \sim \text{Unif}(\{\pm 1\}^d)$ where $q \geq p$ for a suitable value of λ . Our results use the $(2, q)$ -Hypercontractivity Theorem stated below.

Fact 4 ((2, q)-Hypercontractivity Theorem) *For any $f \in \mathcal{F}_{\text{BOOLEAN}}(d)$, we have,*

$$\mathbb{E}_{\text{Unif}(\{\pm 1\}^d)} \left[\text{SMOOTH}_{\sqrt{q-1}}[f](\mathbf{Z})^q \right] \leq \left(\mathbb{E}_{\text{Unif}(\{\pm 1\}^d)} [f^2(\mathbf{Z})] \right)^{\frac{q}{2}}.$$

H.2 Fourier Analysis on the Gauss Space

Consider the functional space $\mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d))$ defined as follows:

$$\mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d)) \stackrel{\text{def}}{=} \left\{ f : \mathbb{R}^d \mapsto \mathbb{R}, \mathbb{E}_{\mathcal{N}(\mathbf{0}, \mathbf{I}_d)} f^2(\mathbf{Z}) < \infty \right\}.$$

The multivariate Hermite polynomials form a complete orthonormal basis for $\mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d))$. These are defined as follows: for any $\mathbf{c} \in \mathbb{N}_0^d$, define,

$$H_{\mathbf{c}}(\mathbf{z}) = \prod_{i=1}^d H_{c_i}(z_i).$$

In the above display, for any $k \in \mathbb{N}_0$, H_k are the usual (univariate) orthonormal Hermite polynomials with the property

$$\mathbb{E}_{\mathcal{N}(0,1)} H_k(Z) H_l(Z) = \begin{cases} 0 & k \neq l \\ 1 & k = l \end{cases}.$$

The orthonormality property is inherited by the multivariate Hermite polynomials and we have,

$$\mathbb{E}_{\mathcal{N}(\mathbf{0}, \mathbf{I}_d)} H_{\mathbf{c}}(\mathbf{Z}) H_{\mathbf{d}}(\mathbf{Z}) = \begin{cases} 0 & \mathbf{c} \neq \mathbf{d} \\ 1 & \mathbf{c} = \mathbf{d} \end{cases}.$$

Since these polynomials form an orthonormal basis of $\mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d))$ any $f \in \mathcal{L}_2(\mathcal{N}(\mathbf{0}, \mathbf{I}_d))$ admits an expansion of the form:

$$f(\mathbf{z}) = \sum_{\mathbf{c} \in (\mathbb{N} \cup \{0\})^d} \hat{f}(\mathbf{c}) H_{\mathbf{c}}(\mathbf{z}).$$

In the above display $\hat{f}(\mathbf{c}) \in \mathbb{R}$ are the Fourier coefficients of f . They satisfy the usual Parseval's relation:

$$\sum_{\mathbf{c} \in \mathbb{N}_0^d} \hat{f}^2(\mathbf{c}) = \mathbb{E}_{\mathcal{N}(\mathbf{0}, \mathbf{I}_d)} f^2(\mathbf{Z}).$$

A particular desirable property of the univariate Hermite polynomials is the following: for any $\mu \in \mathbb{R}, k \in \mathbb{N}_0$ we have,

$$\mathbb{E}_{\mathcal{N}(0,1)} H_k(\mu + Z) = \frac{\mu^k}{\sqrt{k!}}.$$

This implies the following property of multivariate Hermite polynomials which will be particularly useful for us:

Fact 5 For any $\boldsymbol{\mu} \in \mathbb{R}^d$ and any $\mathbf{c} \in \mathbb{N}_0^d$, we have,

$$\mathbb{E}_{\mathcal{N}(\mathbf{0}, \mathbf{I}_d)} H_{\mathbf{c}}(\boldsymbol{\mu} + \mathbf{Z}) = \frac{\boldsymbol{\mu}^{\mathbf{c}}}{\sqrt{\mathbf{c}!}}.$$

In the above display, we are using the following notation:

$$\boldsymbol{\mu}^{\mathbf{c}} \stackrel{\text{def}}{=} \prod_{i=1}^m \mu_i^{c_i}, \quad \mathbf{c}! \stackrel{\text{def}}{=} \prod_{i=1}^m (c_i!).$$

References

- Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor PCA. In *Conference on Learning Theory*, pages 79–104, 2017.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.
- Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor PCA. *arXiv preprint arXiv:1808.00921*, 2018.
- Gerard Ben Arous, Song Mei, Andrea Montanari, and Mihai Nica. The landscape of the spiked tensor model. *Communications on Pure and Applied Mathematics*, 72(11):2282–2330, 2019.
- Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.
- 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.
- 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.
- David Blackwell. Conditional expectation and unbiased sequential estimation. *The Annals of Mathematical Statistics*, pages 105–110, 1947.
- Avrim Blum, Merrick Furst, Jeffrey Jackson, Michael Kearns, Yishay Mansour, and Steven Rudich. Weakly learning DNF and characterizing statistical query learning using Fourier analysis. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 253–262, 1994.
- Matthew Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. *arXiv preprint arXiv:2005.08099*, 2020.
- Matthew Brennan, Guy Bresler, Samuel B Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint arXiv:2009.06107*, 2020.
- Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE, 2017.
- Ilias Diakonikolas, Weihao Kong, and Alistair Stewart. Efficient algorithms and lower bounds for robust linear regression. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2745–2754. SIAM, 2019.

- Vitaly Feldman. Dealing with range anxiety in mean estimation via statistical queries. In *International Conference on Algorithmic Learning Theory*, pages 629–640, 2017a.
- Vitaly Feldman. A general characterization of the statistical query complexity. In *Conference on Learning Theory*, pages 785–830, 2017b.
- Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM (JACM)*, 64(2):8, 2017a.
- Vitaly Feldman, Cristobal Guzman, and Santosh Vempala. Statistical query algorithms for mean vector estimation and stochastic convex optimization. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1265–1277. SIAM, 2017b.
- Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. *SIAM Journal on Computing*, 47(4):1294–1338, 2018.
- 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.
- 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.
- 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.
- Samuel Brink Klevit Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.
- Iain M Johnstone et al. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.
- Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.
- 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.
- Jerry Li, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. On mean estimation for general norms with statistical queries. *arXiv preprint arXiv:1902.02459*, 2019.
- Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.

- Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- C Radhakrishna Rao. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer, 1992.
- Lev Reyzin. Statistical queries and statistical algorithms: Foundations and applications. *arXiv preprint arXiv:2004.00557*, 2020.
- Emile Richard and Andrea Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.
- Joel A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- Zhaoran Wang, Quanquan Gu, and Han Liu. Sharp computational-statistical phase transitions via oracle computational model. *arXiv preprint arXiv:1512.08861*, 2015.
- Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- Anru Zhang and Dong Xia. Tensor SVD: Statistical and computational limits. *IEEE Transactions on Information Theory*, 64(11):7311–7338, 2018.
- Qinqing Zheng and Ryota Tomioka. Interpolating convex and non-convex tensor decompositions via the subspace norm. In *Advances in Neural Information Processing Systems*, pages 3106–3113, 2015.