

Sum-of-squares meets square loss: Fast rates for agnostic tensor completion

Dylan J. Foster Andrej Risteski
 Massachusetts Institute of Technology
 {dylanf, risteski}@mit.edu

Abstract

We study tensor completion in the *agnostic setting*. In the classical tensor completion problem, we receive n entries of an unknown rank- r tensor and wish to exactly complete the remaining entries. In agnostic tensor completion, we make *no assumption* on the rank of the unknown tensor, but attempt to *predict* unknown entries as well as the best rank- r tensor.

For agnostic learning of third-order tensors with the square loss, we give the first polynomial time algorithm that obtains a “fast” (i.e., $O(1/n)$ -type) rate improving over the rate obtained by reduction to matrix completion. Our prediction error rate to compete with the best $d \times d \times d$ tensor of rank- r is $\tilde{O}(r^2 d^{3/2}/n)$. We also obtain an exact oracle inequality that trades off estimation and approximation error.

Our algorithm is based on the degree-six sum-of-squares relaxation of the tensor nuclear norm. The key feature of our analysis is to show that a certain characterization for the subgradient of the tensor nuclear norm can be encoded in the sum-of-squares proof system. This unlocks the standard toolbox for localization of empirical processes under the square loss, and allows us to establish restricted eigenvalue-type guarantees for various tensor regression models, with tensor completion as a special case. The new analysis of the relaxation complements [Barak and Moitra \(2016\)](#), who gave slow rates for agnostic tensor completion, and [Potechin and Steurer \(2017\)](#), who gave exact recovery guarantees for the noiseless setting. Our techniques are user-friendly, and we anticipate that they will find use elsewhere.

1 Introduction

Recovering structured mathematical objects from partial measurements is a fundamental task in machine learning and statistical inference. One important example, which has been a mainstay of modern research in machine learning and high-dimensional statistics, is *matrix completion*. Here, we receive n entries from an unknown $d \times d$ matrix, and the goal is to complete the remaining entries when n is as small as possible. The key structural assumption that enables recovery when $n \ll d^2$ is that the underlying matrix is *low-rank*. A celebrated line of research on matrix completion ([Srebro and Shraibman, 2005](#); [Candès and Recht, 2009](#); [Candès and Tao, 2010](#); [Keshavan et al., 2010](#); [Gross, 2011](#); [Recht, 2011](#)) has culminated in the following guarantee: to *exactly* recover an incoherent rank- r matrix, $n = \tilde{O}(rd)$ uniformly sampled entries suffice.

While low-rank matrix completion has seen successful application across many problem domains (most famously in the context of the *Netflix Problem*), for many tasks it is natural to consider not just pairwise interactions but higher-order interactions, leading to the problem of *tensor completion*. Tensor completion poses significant computational hurdles compared to the matrix case, but in an impressive recent work,

Potechin and Steurer (2017) provide an efficient algorithm based on the sum-of-squares hierarchy that exactly recovers a $d \times d \times d$ tensor of rank- r with incoherent and orthogonal components from $\tilde{O}(rd^{3/2})$ measurements. While this undershoots the optimal statistical rate of $\tilde{O}(rd)$, there is evidence that this is optimal amongst polynomial time algorithms under certain average-case hardness assumptions (Barak and Moitra, 2016).

In real-world applications, the assumption that the underlying tensor is truly low-rank may be too strong, and model misspecification is unavoidable. This is the main thrust of *agnostic learning* (Haussler, 1992; Kearns et al., 1994) which, rather than attempting to recover an unknown model in some class, attempts to *predict* as well as the best model. The aim of this paper is to develop guarantees for *agnostic tensor completion*: predicting as well as the best low rank tensor, even when exact recovery is impossible.

We work in the following *agnostic tensor regression* model, which captures tensor completion as a special case: we receive examples $(X_1, Y_1), \dots, (X_n, Y_n)$ i.i.d. from an unknown distribution \mathcal{D} , where each instance X_t is a $d \times d \times d$ tensor and Y_t is a real-valued response. Letting $\langle \cdot, \cdot \rangle$ denote the usual inner product, we measure predictive performance of a given $d \times d \times d$ tensor T via its *square loss risk* $L_{\mathcal{D}}(T) = \mathbb{E}_{(X, Y) \sim \mathcal{D}}(\langle T, X \rangle - Y)^2$. Our goal is to use the samples to produce a predictor \widehat{T}_n that enjoy low *excess risk*:

$$L_{\mathcal{D}}(\widehat{T}_n) - \inf_{T: \text{rank-}r} L_{\mathcal{D}}(T) \leq \varepsilon(n, r, d), \quad (1)$$

where the bound $\varepsilon(n, r, d)$ converges to zero as $n \rightarrow \infty$. When the observations X are uniformly distributed indicators (i.e. $X = e_i \otimes e_j \otimes e_k$, where (i, j, k) is uniform) this recovers the usual measurement model for tensor completion. If the model is *well-specified* in the sense that $Y = \langle X, T^* \rangle + \xi$, where T^* is a low-rank tensor and $\mathbb{E}[\xi | X] = 0$, then low excess risk implies approximate recovery of T^* . In general, the guarantee (1) is interesting because it implies non-trivial predictive performance even in the presence of severe model misspecification.

In the matrix case, agnostic excess risk guarantees of the type in (1) were characterized by Koltchinskii et al. (2011). There it was shown that to compete with the best rank- r matrix with bounded entries, empirical risk minimization with nuclear norm penalization obtains a *fast rate* of the form $\varepsilon(n, r, d) = O(rd \log d/n)$, and also showed that this is optimal.¹ The nomenclature fast rate is intended to contrast the $1/n$ dependency on n with *slow rates*, which have a $1/\sqrt{n}$ dependency on n , and which are typically much simpler to obtain (see Gaiffas and Lecué (2011) for slow rates in the matrix setting).

The results of Koltchinskii et al. (2011) leverage strong understanding of the (matrix) nuclear norm, namely matrix concentration and decomposability/subgradient properties. In this paper, we tackle the following questions:

- Can we give similar guarantees for agnostic *tensor* completion?
- Can we obtain fast rates while at the same time relaxing the strong statistical assumptions (low rank observations, incoherence) needed for exact recovery?
- What are the best rates we can obtain subject to employing a polynomial-time algorithm?

¹Interestingly, their results also show that incoherence—usually taken as necessary in matrix completion for positive results—is not necessary to obtain prediction error bounds; boundedness suffices.

1.1 Our contributions

Our main result is to give the first polynomial time algorithm with a fast $O(1/n)$ -type rate for agnostic completion of third order tensors that improves over the rate obtained by the natural reduction to matrix completion. Our main theorem gives excess risk bounds relative to low-rank *orthogonal* tensors, i.e. tensors of the form

$$T = \sum_{i=1}^r \lambda_i \cdot u_i \otimes v_i \otimes w_i, \quad (2)$$

where $\|u_i\|_2 = \|v_i\|_2 = \|w_i\|_2 = 1$ and $\{u_i\}$ are orthogonal, as are $\{v_i\}$ and $\{w_i\}$. The result is as follows.

Theorem 1 (informal). *There is a convex set of $d \times d \times d$ tensors \mathcal{T} derived from a sum-of-squares relaxation of the tensor nuclear norm, for which the empirical risk minimizer $\widehat{T}_n := \arg \min_{T \in \mathcal{T}} \frac{1}{n} \sum_{t=1}^n (\langle T, X_t \rangle - Y_t)^2$ can be computed in polynomial time, and guarantees that with probability at least $1 - o(1)$,*

$$L_{\mathcal{D}}(\widehat{T}_n) - \inf_{T: \text{rank-}r, \text{ orthogonal}} L_{\mathcal{D}}(T) \leq \widetilde{O}\left(\frac{r^2 d^{3/2}}{n}\right). \quad (3)$$

The guarantee applies to both random indicator measurements (tensor completion) and gaussian measurements (tensor compressed sensing).

The full version of the theorem is stated in [Section 3](#). The result be thought of as a generalization of the agnostic matrix completion results of [Koltchinskii et al. \(2011\)](#) to higher-order tensors. We also achieve a more general *exact oracle inequality* that trades off approximation and estimation error. This takes the form

$$L_{\mathcal{D}}(\widehat{T}_n) \leq \inf_{T: \text{orthogonal}} \left\{ L_{\mathcal{D}}(T) + \widetilde{O}\left(\frac{r^2(T) d^{3/2}}{n}\right) \right\},$$

where $r(T)$ denotes the rank of T .²

Algorithm. Our algorithm is based on the sum-of-squares (SoS) hierarchy of convex relaxations ([Shor, 1987; Parrilo, 2000; Lasserre, 2001](#)), applied to the tensor nuclear norm. For $k \in \mathbb{N}$, the sum-of-squares hierarchy defines a sequence of outer convex relaxations $\|\cdot\|_{\text{nuc}_k}$ that give an increasingly tight approximation to the (hard to approximate ([Hillar and Lim, 2013](#))) tensor nuclear norm $\|\cdot\|_{\text{nuc}}$ as k is increased. The SoS nuclear norm was previously used in the work of [Barak and Moitra \(2016\)](#), who gave $O(1/\sqrt{n})$ rates for tensor completion in the agnostic model with absolute loss. Their main contribution was to show how certain spectral bounds arising in 3-SAT refutation ([Coja-Oghlan et al., 2004](#)) bound the Rademacher complexity for the unit ball of the degree-six SoS nuclear norm, thereby controlling the usual empirical process via uniform convergence. Our guarantees are based on empirical risk minimization over the (scaled) ball in this norm, specifically the degree-six relaxation $\|\cdot\|_{\text{nuc}_6}$, and our analysis builds on their Rademacher complexity bounds.

Restricted eigenvalue and subgradient lemmas. While the $O(1/\sqrt{n})$ -type rates provided in [Barak and Moitra \(2016\)](#) are optimal (in terms of n dependence) for generic Lipschitz losses, it is not immediately obvious whether their result can be used to provide $O(1/n)$ -type fast rates for strongly convex losses like the square loss. It has long been recognized that to obtain fast rates for prediction with strongly convex losses, more refined control of the empirical process is necessary. In particular, it is well-known that *local* Rademacher complexities and related fixed point complexities characterize the rates for empirical risk minimization in a data-dependent fashion. To bound such localized complexities under convex relaxations, the typical approach is to establish a restricted eigenvalue property for the empirical design matrix ([Negahban et al., 2012](#);

²Here the word “exact” refers to the fact that the leading constant in front of the loss on the right-hand-side is 1; such guarantees generally do not easily follow from the usual machinery used to analyze well-specified models.

Bartlett et al., 2012; Lecué and Mendelson, 2017, 2018). Establishing such guarantees for regularizers such as the ℓ_1 norm, nuclear norm, and so on is usually done by appealing to properties of the subgradient of the norm and proving that the norm (approximately) decomposes across certain subspaces (Negahban et al., 2012). The main tool we establish here, which allows us to unlock the full power of localized complexities and establish rates fast rates, is a new guarantee of this type for the SoS nuclear norm. Informally, we show:

Theorem 2 (informal). *Let T^* be a fixed orthogonal tensor, and let $\mathcal{T} = \{T : \|T\|_{\widetilde{\text{nuc}}_6} \leq \|T^*\|_{\widetilde{\text{nuc}}_6}\}$. Then:*

$$\|T - T^*\|_{\widetilde{\text{nuc}}_4} \leq O(r(T^*)) \cdot \|T - T^*\|_F \quad \forall T \in \mathcal{T}.$$

This theorem is a consequence of a more general result we prove in Section 2, which gives a characterization of the subgradient of $\|\cdot\|_{\widetilde{\text{nuc}}_4}$ at any orthogonal tensor. The basic idea is to show that a characterization for the subgradient of the tensor nuclear norm given in Yuan and Zhang (2016) can be captured in the sum-of-squares proof system. The final result (Theorem 4) is fairly user-friendly, and packages the complexity of SoS into a self-contained statement about geometric properties of the norm $\|\cdot\|_{\widetilde{\text{nuc}}_4}$. We hope that this will find use in other applications. As one concrete example, in Section 3 we show how the subgradient characterization can also be used to give agnostic fast rates for the problem of *low-rank tensor sensing* with gaussian measurements. Here we also obtain $\widetilde{O}(r^2 d^{3/2}/n)$ -type fast rates.

Lower bounds. Lastly, we prove that our results can't be strengthened significantly while still obtaining computationally efficient algorithms. It is straightforward to show that there is an *inefficient* predictor that obtains $O(rd/n)$ excess risk, whereas the dimension scaling in Theorem 1 is $O(d^{3/2})$. This scaling is shared by the other results based on the SoS nuclear norm (Barak and Moitra, 2016; Potechin and Steurer, 2017), and Barak and Moitra (2016) show that finding relaxations for which the Rademacher complexity grows as $o(d^{3/2})$ is at least as hard as refuting random instances of 3-XOR. In Section 4 we give a computational lower bound for agnostic learning that shows that obtaining square loss excess risk scaling as $o(d^{3/2})$ is at least as hard as a certain distinguishing problem for random 3-XOR, sometimes called learning sparse parities with noise.

1.2 Related work

Early algorithms for computationally efficient tensor completion relied on *unfolding*: reshaping the tensor into a matrix and applying a matrix completion algorithm (Tomioka et al., 2011; Tomioka and Suzuki, 2013; Romera-Paredes and Pontil, 2013; Mu et al., 2014; Jain and Oh, 2014). This approach yields suboptimal results for third-order tensors and other odd-order tensors. For example, for a third-order tensor in $(\mathbb{R}^d)^{\otimes 3}$, the most “balanced” unfolding of the tensor is a $d \times d^2$ matrix, and so directly reducing to an algorithm for agnostic matrix completion (e.g. Koltchinskii et al. (2011)) would yield suboptimal $O(d^2)$ -type sample complexity.

Recent results of Montanari and Sun (2016), Potechin and Steurer (2017), and Xia et al. (2017) all give sub- $O(d^2)$ type rates, but apply only to noiseless or well-specified models, and do not obviously extend to the agnostic setting. Potechin and Steurer (2017) give *exact completion* in the noiseless case after $\widetilde{O}(rd^{3/2})$ entries are observed. Montanari and Sun (2016) show that a refined spectral approach based on unfolding can obtain sub- $O(d^2)$ rates for prediction error in the noiseless setting, but they obtain a rate of $O(1/n^{2/3})$ that falls short of the $O(1/n)$ -type rate we provide. Finally, Xia et al. (2017) recently showed that an algorithm based on power iteration provides $O(d/n)$ -type rates once $n = \Omega(d^{3/2})$, but their result only applies to well-specified models, and it seems unlikely that this algorithmic approach succeeds in the fully agnostic setting.

Our results build on the seminal work of Barak and Moitra (2016) and Potechin and Steurer (2017), both of which use sum-of-squares to give $O(d^{3/2})$ -type guarantees that improve on unfolding. The former was the first paper to study the agnostic setting, and gave slow excess risk guarantees for the absolute loss (i.e., rates growing as $\frac{1}{\sqrt{n}}$). Technically, our results build on their Rademacher complexity bounds for the SoS norms (Barak and Moitra, 2016), as well as spectral bounds from Hopkins et al. (2015). Obtaining fast rates, however, requires developing new technical tools and necessitates that we control the subgradient of the SoS norms. Our analysis here builds on ideas used to construct dual certificates for tensor completion in Potechin and Steurer (2017).

Lastly, we mention that various recent works have begun to explore that power of SoS in other agnostic learning settings. Notably, Klivans et al. (2018) provide square loss risk bounds for SoS algorithms for robust regression. To the best of our knowledge our work is the first to provide *fast rates* for SoS algorithms in any agnostic setting.

1.3 Preliminaries

We let $\|\cdot\|_p$ denote the ℓ_p norm, i.e. if $x \in \mathbb{R}^d$ is a vector then $\|x\|_p = (\sum_{i=1}^d |x_i|^p)^{1/p}$. For a matrix A we let $\|A\|_{\text{op}}$ denote the operator norm/spectral norm and let $\|A\|_{\text{nuc}}$ denote the nuclear norm. For matrices or tensors we let $\|\cdot\|_F$ denote the element-wise ℓ_2 norm. For any norm $\|\cdot\|$ we let $\|\cdot\|_*$ denote the dual.

We use \widetilde{O} to suppress factors logarithmic in d , r , and $1/\delta$, where δ is the failure probability.

Tensor notation. The outer product between two vectors $u \in \mathbb{R}^{d_1}$ and $v \in \mathbb{R}^{d_2}$ is denoted $u \otimes v$, and belongs to the space $\mathbb{R}^{d_1} \otimes \mathbb{R}^{d_2}$. For a given vector u , we write $u^{\otimes k} = u \otimes \cdots \otimes u$ (k times), and likewise define $(\mathbb{R}^d)^{\otimes k} = \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d$. This paper develops algorithms for completion of 3-tensors in $(\mathbb{R}^d)^{\otimes 3}$, which we frequently identify with elements of $\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d$. In more detail, for a tensor $T \in (\mathbb{R}^d)^{\otimes 3}$, we let $T_{i,j,k}$ be such that $T = \sum_{i,j,k} T_{i,j,k} \cdot e_i \otimes e_j \otimes e_k$. For a pair of matrices A, B , we let $A \otimes B$ denote the Kronecker product, which obeys the relation $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$. Given matrices A_1, A_2, A_3 and tensor $T = \sum_i u_i \otimes v_i \otimes w_i$, we define $(A_1 \otimes A_2 \otimes A_3)T = \sum_i (A_1 u_i) \otimes (A_2 v_i) \otimes (A_3 w_i)$. Whenever T is an orthogonal tensor of the form (2), we let $r(T)$ denote the rank.

2 Subgradient of the sum-of-squares nuclear norm

2.1 Tensor nuclear norm and sum-of-squares relaxation

In the classical results on matrix completion (Candès and Recht, 2009; Candès and Tao, 2010; Recht, 2011; Koltchinskii et al., 2011), a central object is the *nuclear* norm, which arises as a convex relaxation of the rank. A natural candidate to develop efficient algorithms for tensor completion is the *tensor nuclear norm* (e.g. Hillar and Lim (2013); Friedland and Lim (2018)), which may be defined via

$$\|T\|_{\text{nuc}} = \inf \left\{ \sum_{i=1}^r |\lambda_i| : W = \sum_{i=1}^r \lambda_i u_i \otimes v_i \otimes w_i, \|u_i\|_2 = \|v_i\|_2 = \|w_i\|_2 = 1, r \in \mathbb{N} \right\}, \quad (4)$$

and whose dual is the *injective tensor norm* $\|X\|_{\text{inj}} = \sup_{\|x\|_2 = \|y\|_2 = \|z\|_2 = 1} \langle X, x \otimes y \otimes z \rangle$. Unfortunately, the optimization problem here—maximizing a degree-three polynomial over the sphere—is intractable in general.

The approach we take, following Barak and Moitra (2016) and Potechin and Steurer (2017), is to employ the sum-of-squares hierarchy of convex relaxations (Shor, 1987; Parrilo, 2000; Lasserre, 2001) which provides an increasingly tight sequence of relaxations of the optimization problem in (4). To describe the relaxations, we require the notion of a *pseudodistribution*.

Definition 1 (Pseudodistribution (Barak and Steurer (2016))). *Let $\mu : \mathbb{R}^d \rightarrow \mathbb{R}$ be a finitely supported function and let $\widetilde{\mathbb{E}}_\mu f = \sum_{x \in \text{supp}(\mu)} \mu(x) f(x)$. μ is said to be a degree- k pseudodistribution if $\widetilde{\mathbb{E}}_\mu 1 = 1$ and $\widetilde{\mathbb{E}}_\mu f^2 \geq 0$ for all polynomials f of degree at most $k/2$.*

Given a degree- s pseudodistribution μ and system of polynomial inequalities $\mathcal{A} = \{f_1 \geq 0, \dots, f_m \geq 0\} \cup \{g_1 = 0, \dots, g_{m'} = 0\}$, we write $\mu \models \mathcal{A}$ if for all $S \subseteq [m]$ and all sum-of-squares polynomials h such that $\deg h + \sum_{i \in S} \deg f_i \leq s$,

$$\widetilde{\mathbb{E}}_\mu \left[h \prod_{i \in S} f_i \right] \geq 0, \quad (5)$$

and $\widetilde{\mathbb{E}}_\mu [g_i q] = 0$ for all $i \in [m']$ and all polynomials q such that $\deg(g_i q) \leq \ell$.

With the pseudodistribution formalism, we define the degree- k SoS injective norm as follows:³

$$\|X\|_{\widetilde{\mathbb{E}}_k} = \sup_{\substack{\mu \text{ degree-}k \\ \widetilde{\mathbb{E}}_\mu \|x\|_2^2 = \widetilde{\mathbb{E}}_\mu \|y\|_2^2 = \widetilde{\mathbb{E}}_\mu \|z\|_2^2 = 1}} \langle X, \widetilde{\mathbb{E}}_\mu [x \otimes y \otimes z] \rangle. \quad (6)$$

The degree- k SoS nuclear norm is simply defined as the dual: $\|T\|_{\widetilde{\mathbb{E}}_k} := \sup_{X: \|X\|_{\widetilde{\mathbb{E}}_k} \leq 1} \langle X, T \rangle$. It can equivalently be expressed by defining⁴

$$\mathcal{K}_k = \{T \in \mathbb{R}^{d \times d \times d} \mid \exists \mu \text{ degree-}k : \widetilde{\mathbb{E}}_\mu [x \otimes y \otimes z] = T, \widetilde{\mathbb{E}}_\mu \|x\|_2^2 = \widetilde{\mathbb{E}}_\mu \|y\|_2^2 = \widetilde{\mathbb{E}}_\mu \|z\|_2^2 = 1\}, \quad (7)$$

and then

$$\|T\|_{\widetilde{\mathbb{E}}_k} = \inf \{\alpha \mid T \in \mathcal{K}_k / \alpha\}. \quad (8)$$

The SoS nuclear norm and injective norm can be evaluated in $d^{O(k)}$ time (Grötschel et al., 1981; Barak and Steurer, 2016). Moreover, the norms obey the ordering $\|T\|_{\text{nuc}} \geq \dots \geq \|T\|_{\widetilde{\mathbb{E}}_k} \geq \dots \geq \|T\|_{\widetilde{\mathbb{E}}_2} \geq \|T\|_F$, and likewise $\|X\|_{\text{inj}} \leq \dots \leq \|X\|_{\widetilde{\mathbb{E}}_k} \leq \dots \leq \|X\|_{\widetilde{\mathbb{E}}_2} \leq \|X\|_F$.

2.2 Subgradient and norm compatibility

Our algorithms are based on empirical risk minimization with SoS nuclear norm constraints, i.e., algorithms that minimize the empirical loss over the set $\mathcal{T} = \{T \in (\mathbb{R}^d)^{\otimes 3} \mid \|T\|_{\widetilde{\mathbb{E}}_k} \leq \tau\}$ for appropriate choice of τ and k . The technical challenge to analyzing this type of relaxation is that even if measurements are realized by a rank- r tensor, there is nothing that guarantees a-priori that the tensor \widehat{T}_n output by the algorithm is itself low rank. Letting T be a rank- r orthogonal tensor, our main technical result here shows that if $\tau = \|T\|_{\widetilde{\mathbb{E}}_k}$, then for all elements of $T' \in \mathcal{T}$, the error $\Delta = T' - T$ is “approximately low rank” in a sense that suffices to guarantee good generalization performance.

Theorem 3 (Formal version of Theorem 2). *Let $k \geq 4$. Let $T \in (\mathbb{R}^d)^{\otimes 3}$ be an orthogonal rank- r tensor, and let $T' \in (\mathbb{R}^d)^{\otimes 3}$ be an arbitrary tensor with $\Delta := T' - T$. If $\|T'\|_{\widetilde{\mathbb{E}}_{k+2}} \leq \|T\|_{\widetilde{\mathbb{E}}_{k+2}}$, then*

$$\|\Delta\|_{\widetilde{\mathbb{E}}_k} \leq 68r \cdot \|\Delta\|_F. \quad (9)$$

³There are minor technical differences (e.g., scaling) between the SoS injective/nuclear norm definitions we use and those of Barak and Moitra (2016); Potechin and Steurer (2017).

⁴This equivalence is proven in Appendix A for completeness.

In the analysis of nuclear norm regularization for matrix completion—and more broadly, throughout high-dimensional statistics—the key tool used to establish guarantees along the lines of (9) is a characterization for the *subgradient* for the nuclear norm, and the related notion of *decomposability* (Negahban et al., 2012; Negahban and Wainwright, 2012). It is known classically (Watson, 1992) that for any matrix W with singular value decomposition $W = U\Sigma V^\top$,

$$\partial\|W\|_{\text{nuc}} = \{UV^\top + X \mid U^\top X = XV = 0, \|X\|_{\text{op}} \leq 1\}. \quad (10)$$

Our approach in the remainder of this section is to establish a similar result for the subgradient of the SoS nuclear norm $\|\cdot\|_{\widetilde{\text{nuc}}_k}$ at any orthogonal tensor T . From here [Theorem 3](#) will quickly follow.

As a first step, we need to define certain subspaces and projection operators associated with T .

Subspaces For the remainder of the section we let T be a rank- r orthogonal tensor as in (2). Define $\mathbb{U} = \text{span}\{u_i\}$, $\mathbb{V} = \text{span}\{v_i\}$, and $\mathbb{W} = \text{span}\{w_i\}$, and note that each subspace has dimension at most r . Let $\mathcal{P}_{\mathbb{U}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\mathcal{P}_{\mathbb{V}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\mathcal{P}_{\mathbb{W}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be orthogonal projections onto these subspaces and $\mathcal{P}_{\mathbb{U}^\perp}$, $\mathcal{P}_{\mathbb{V}^\perp}$ and $\mathcal{P}_{\mathbb{W}^\perp}$ be the projections onto the respective orthogonal complements. We define projection operators from $(\mathbb{R}^d)^{\otimes 3}$ to $(\mathbb{R}^d)^{\otimes 3}$ for all 2^3 combinations of subspaces:⁵

$$\begin{aligned} \mathcal{Q}_{T\parallel}^0 &= \mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}}, & \mathcal{Q}_{T^\perp}^0 &= \mathcal{P}_{\mathbb{U}^\perp} \otimes \mathcal{P}_{\mathbb{V}^\perp} \otimes \mathcal{P}_{\mathbb{W}^\perp}, \\ \mathcal{Q}_{T\parallel}^1 &= \mathcal{P}_{\mathbb{U}^\perp} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}}, & \mathcal{Q}_{T^\perp}^1 &= \mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}^\perp} \otimes \mathcal{P}_{\mathbb{W}^\perp}, \\ \mathcal{Q}_{T\parallel}^2 &= \mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}^\perp} \otimes \mathcal{P}_{\mathbb{W}}, & \mathcal{Q}_{T^\perp}^2 &= \mathcal{P}_{\mathbb{U}^\perp} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}^\perp}, \\ \mathcal{Q}_{T\parallel}^3 &= \mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}^\perp}, & \mathcal{Q}_{T^\perp}^3 &= \mathcal{P}_{\mathbb{U}^\perp} \otimes \mathcal{P}_{\mathbb{V}^\perp} \otimes \mathcal{P}_{\mathbb{W}}. \end{aligned} \quad (11)$$

Lastly, we define two subspaces that play a central role in our analysis:

$$\mathcal{Q}_{T\parallel} = \mathcal{Q}_{T\parallel}^0 + \mathcal{Q}_{T\parallel}^1 + \mathcal{Q}_{T\parallel}^2 + \mathcal{Q}_{T\parallel}^3, \quad \text{and} \quad \mathcal{Q}_{T^\perp} = \mathcal{Q}_{T^\perp}^0 + \mathcal{Q}_{T^\perp}^1 + \mathcal{Q}_{T^\perp}^2 + \mathcal{Q}_{T^\perp}^3. \quad (12)$$

One can verify via multilinearity that $X = \mathcal{Q}_{T\parallel}(X) + \mathcal{Q}_{T^\perp}(X)$ for all $X \in (\mathbb{R}^d)^{\otimes 3}$, and that any tensor in the range of $\mathcal{Q}_{T\parallel}$ spans at most r dimensions along at least two modes. We can now state our main theorem for the subgradient.

Theorem 4 (Subgradient of SoS nuclear norm). *Let $k \geq 4$, and let $T = \sum_{i=1}^r \lambda_i \cdot u_i \otimes v_i \otimes w_i$ be an orthogonal rank- r tensor. Define $X^* = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$. Then for all $X \in (\mathbb{R}^d)^{\otimes 3}$ with $\|X\|_{\widetilde{\text{nuc}}_k} \leq 1/64$, and for all $T' \in (\mathbb{R}^d)^{\otimes 3}$, it holds that*

$$\|T'\|_{\widetilde{\text{nuc}}_{k+2}} \geq \|T\|_{\widetilde{\text{nuc}}_{k+2}} + \langle X^* + \mathcal{Q}_{T^\perp}(X), T' - T \rangle. \quad (13)$$

In other words, [Theorem 4](#) states that

$$\{X^* + \mathcal{Q}_{T^\perp}(X) \mid \|X\|_{\widetilde{\text{nuc}}_k} \leq 1/64\} \subset \partial\|T\|_{\widetilde{\text{nuc}}_{k+2}},$$

which we may view as a generalization of the matrix subgradient characterization (10). [Yuan and Zhang \(2016\)](#) proved a similar result for the (exact) tensor nuclear norm. Our proof of [Theorem 4](#) shows that the essence of their proof can be captured by a low-degree sum-of-squares proof. It builds on the approach introduced in [Potechin and Steurer \(2017\)](#) to provide dual certificates for exact tensor completion.

⁵Following the convention in [Section 1.3](#), if $X = \sum_i a_i \otimes b_i \otimes c_i$, then $(\mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}})X = \sum_i (\mathcal{P}_{\mathbb{U}}a_i) \otimes (\mathcal{P}_{\mathbb{V}}b_i) \otimes (\mathcal{P}_{\mathbb{W}}c_i)$. It is also useful to note that for any x, y, z we have $\langle (\mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}}), x \otimes y \otimes z \rangle = \langle X, (\mathcal{P}_{\mathbb{U}}x) \otimes (\mathcal{P}_{\mathbb{V}}y) \otimes (\mathcal{P}_{\mathbb{W}}z) \rangle$.

With the subgradient lemma in hand, the path to the “approximately low rank” result of [Theorem 3](#) is clear. Suppose that $\|T'\|_{\widetilde{\text{nuc}}_{k+2}} \leq \|T\|_{\widetilde{\text{nuc}}_{k+2}}$, and let $\Delta = T' - T$. By appropriately choosing the dual tensor X in [\(13\)](#), we can show that

$$\|\mathcal{Q}_{T^\perp}(\Delta)\|_{\widetilde{\text{nuc}}_k} \leq 64 \cdot \|\mathcal{Q}_{T^\parallel}^0(\Delta)\|_{\widetilde{\text{nuc}}_k}. \quad (14)$$

which implies that $\|\Delta\|_{\widetilde{\text{nuc}}_k} \leq 65 \|\mathcal{Q}_{T^\parallel}(\Delta)\|_{\widetilde{\text{nuc}}_k}$. The final result follows because $\|\mathcal{Q}_{T^\parallel}(\Delta)\|_{\widetilde{\text{nuc}}_k} \lesssim r \|\Delta\|_F$, which is a consequence of the earlier remark that all the projections used to define $\mathcal{Q}_{T^\parallel}$ in [\(12\)](#) project into r dimensions along at least two modes. This full argument is in [Appendix B](#).

3 Agnostic tensor completion

We now state our main learning results, which use the SoS nuclear norm to give efficient algorithms with fast rates for agnostic tensor completion and tensor sensing. For both results we receive observations $(X_1, Y_1), \dots, (X_n, Y_n)$ i.i.d. according to an unknown distribution \mathcal{D} , where $(X_t, Y_t) \in (\mathbb{R}^d)^{\otimes 3} \times \mathbb{R}$, and the goal is to obtain low square loss excess risk in the sense of equation [\(1\)](#). We let $\widehat{\mathbb{E}}_n$ denote the empirical expectation, which is uniform over the examples $\{(X_t, Y_t)\}_{t=1}^n$.

3.1 Tensor completion

In the tensor completion model we take observations X_t to be of the form $X_t = e_{i_t} \otimes e_{j_t} \otimes e_{k_t}$, where $(i_t, j_t, k_t) \in [d]^3$ is selected uniformly at random.⁶ In the noiseless or well-specified setting, this corresponds to observing a single entry of an unknown tensor, but we make no assumption on the responses Y_t other than boundedness. The main theorem is as follows.

Theorem 5 (Formal version of [Theorem 1](#)). *Let $\tau > 0$ be fixed. Suppose that $|Y| \leq R$ almost surely, and let \widehat{T}_n be the empirical risk minimizer over the tensor class $\mathcal{T} = \{T \in (\mathbb{R}^d)^{\otimes 3} \mid \|T\|_{\widetilde{\text{nuc}}_6} \leq \tau, \|T\|_\infty \leq R\}$. Then for all $n \leq d^3$, with probability at least $1 - \delta$,*

$$L_{\mathcal{D}}(\widehat{T}_n) - L_{\mathcal{D}}(T^*) \leq O\left(\frac{R^2 r^2 (T^*) d^{3/2} \log^6 d}{n} + \frac{R^2 \log(1/\delta)}{n}\right) \quad (15)$$

for all orthogonal tensors $T^* \in (\mathbb{R}^d)^{\otimes 3}$ with $\|T^*\|_{\widetilde{\text{nuc}}_6} = \tau$ and $\|T^*\|_\infty \leq R$.

Let us spend a moment interpreting the theorem. First, let $T^* = \arg \min_{T: \text{rank-}r} L_{\mathcal{D}}(T)$. Then, by setting $\tau = \|T^*\|_{\widetilde{\text{nuc}}_6}$, we are guaranteed that with probability at least $1 - \delta$,

$$L_{\mathcal{D}}(\widehat{T}_n) - \inf_{T: \text{rank-}r} L_{\mathcal{D}}(T) \leq O\left(\frac{R^2 r^2 d^{3/2} \log^6 d}{n} + \frac{R^2 \log(1/\delta)}{n}\right). \quad (16)$$

More generally, [\(15\)](#) implies an *exact oracle inequality* ([Koltchinskii et al., 2011](#); [Gaiffas and Lecué, 2011](#)): With probability at least $1 - \delta$, we have

$$L_{\mathcal{D}}(\widehat{T}_n) \leq \inf_{T: \|T\|_{\widetilde{\text{nuc}}_6} = \tau} \left\{ L_{\mathcal{D}}(T) + O\left(\frac{R^2 r^2 (T) d^{3/2} \log^6 d}{n} + \frac{R^2 \log(1/\delta)}{n}\right) \right\}.$$

Let us compare the result in detail with [Barak and Moitra \(2016\)](#), which is the only other polynomial time agnostic tensor completion result with sub- $O(d^2)$ sample complexity. For general noise distributions, their

⁶Note that we sample entries with replacement, whereas related works use without-replacement sampling ([Barak and Moitra, 2016](#); [Potechin and Steurer, 2017](#)).

analysis gives an excess risk bound that scales as $\tilde{O}\left(\sqrt{\frac{r^2 d^{3/2}}{n}}\right)$. The bound in (15) matches this dependence on all the parameters, but is squared, and is thus always tighter. The result has excess risk against arbitrary tensors, however, while our bound requires orthogonality of the benchmark. We do not know whether this restriction can be removed.

Interestingly, while [Barak and Moitra \(2016\)](#) give excess risk bounds against incoherent tensors, we do not require incoherence. This is because we control the complexity of the benchmark through the ℓ_∞ norm of the entries rather than through the Frobenius norm; this parallels the situation in the matrix setting ([Koltchinskii et al., 2011](#); [Gaiffas and Lecué, 2011](#)). Applying the spectral bounds of [Barak and Moitra \(2016\)](#) without incoherence requires slightly tightening the analysis.

Lastly, we remark that the guarantee (16) requires setting the parameter τ based on the norm of the unknown benchmark T^* . It is likely that this can be relaxed by appealing to penalized empirical risk minimization rather than empirical risk minimization as in [Koltchinskii et al. \(2011\)](#), but we leave this for future work.

3.2 Tensor sensing

In this setting we give agnostic learning guarantees for a setting we call *tensor sensing*, which generalizes the matrix compressed sensing setup studied in [Negahban and Wainwright \(2011\)](#). We assume that observations $X \in (\mathbb{R}^d)^{\otimes 3}$ have independent entries from $\mathcal{N}(0; 1)$ and—as in the tensor completion setting—allow $Y \in \mathbb{R}$ to be arbitrary. For each tensor T , define $R(T)$ to be the smallest almost-sure bound on $|\langle T, X \rangle - Y|$. As in the tensor completion setup, the main result is a fast rate with $\tilde{O}\left(\frac{r^2 d^{3/2}}{n}\right)$ -type scaling.

Theorem 6. *Let $\tau > 0$ be fixed, and let \widehat{T}_n be the empirical risk minimizer over the tensor class $\mathcal{T} = \{T \in (\mathbb{R}^d)^{\otimes 3} \mid \|T\|_{\text{Hilb}_6} \leq \tau\}$. Then with probability at least $1 - \delta$,*

$$L_{\mathcal{D}}(\widehat{T}_n) - L_{\mathcal{D}}(T^*) \leq O\left(\frac{R^2(T^*) r^2(T^*) d^{3/2} \log^3(d/\delta)}{n}\right). \quad (17)$$

for all orthogonal tensors $T^* \in (\mathbb{R}^d)^{\otimes 3}$ for which $n = \Omega(r^2(T^*) d^{3/2} \log^{1/2} d + \log(1/\delta))$ and $\|T^*\|_{\text{Hilb}_6} = \tau$.

Overview of analysis We now sketch how the subgradient theorem can be combined with empirical process arguments to prove [Theorem 5](#) and [Theorem 6](#). We follow a generic recipe given in [Appendix C.1](#)—specifically, [Theorem 8](#)—which shows that to control the generalization error of empirical risk minimization, it suffices to bound a certain “offset” or “shifted” empirical process. For any fixed benchmark T^* , the excess risk relative to T^* is bounded as

$$L_{\mathcal{D}}(\widehat{T}_n) - L_{\mathcal{D}}(T^*) \leq \sup_{\Delta \in \mathcal{T} - T^*} \{(\mathbb{E} - \widehat{\mathbb{E}}_n)[2\langle \Delta, X \rangle (\langle T^*, X \rangle - Y)] + \mathbb{E}\langle \Delta, X \rangle^2 - 2\widehat{\mathbb{E}}_n\langle \Delta, X \rangle^2\}. \quad (18)$$

The offset process on the right-hand side was used to obtain high-probability fast rates for misspecified models by [Liang et al. \(2015\)](#), and its analysis is closely related to that of [Lecué and Mendelson \(2013\)](#); [Mendelson \(2014\)](#). To bound the process, it suffices to establish a type of lower isometry/restricted eigenvalue property, which we state here for the case of tensor regression: let $\mathcal{X}_n : (\mathbb{R}^d)^{\otimes 3} \rightarrow \mathbb{R}^n$ be the data operator, which maps any tensor T to the sequence $\langle T, X_1 \rangle, \dots, \langle T, X_n \rangle$, and let $\Sigma = \mathbb{E}_X[XX^\top] \in \mathbb{R}^{d^3 \times d^3}$ be the covariance matrix for the vectorized measurements. Then it suffices to show that with high probability, the following *restricted eigenvalue* bound holds:

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq c \|\Sigma^{1/2} \Delta\|_F \quad \forall \Delta \in \mathcal{T} - T^*,$$

where $c > 1/\sqrt{2}$ is a sufficiently large constant.

Our starting point to establish the guarantee is to borrow a bound from [Hopkins et al. \(2015\)](#), which states that $\mathbb{E}\|X\|_{\widetilde{\text{inj}}_4} = O(d^{3/4} \log^{1/4} d)$ under gaussian measurements, and suffices to bound the Rademacher complexity of our tensor class. Using this bound in conjunction with standard gaussian concentration arguments and the “peeling” method (e.g. [\(Negahban and Wainwright, 2011\)](#)), we prove [Theorem 10](#), which states that with high probability,

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq 0.79 \cdot \|\Delta\|_F - \frac{Cd^{3/4} \log^{1/4} d}{\sqrt{n}} \cdot \|\Delta\|_{\widetilde{\text{nuc}}_4} \quad \forall \Delta \in (\mathbb{R}^d)^{\otimes 3}.$$

Combined with [Theorem 3](#), which asserts that all $\Delta \in \mathcal{T} - T^*$ have $\|\Delta\|_{\widetilde{\text{nuc}}_4} \leq O(r(T^*)) \cdot \|\Delta\|_2$, we have the following consequence: once $n = \Omega(r^2(T^*)d^{3/2} \log^{1/2} d)$, with high probability,

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq (0.79 - o(1)) \cdot \|\Delta\|_F \quad \forall \Delta \in \mathcal{T} - T^*.$$

To establish the analogous bound in the tensor completion model we use the SoS Rademacher bound from [Barak and Moitra \(2016\)](#), but utilize somewhat different concentration arguments. Indeed, due to the sparse nature of the measurement distribution one cannot hope to exactly establish the restricted eigenvalue property for \mathcal{X}_n , and must instead show that it holds up to a small additive error.

4 Computational lower bounds

In the rank-one case, the excess risk bound of [Theorem 5](#) scales as $\tilde{O}(d^{3/2}/n)$, while the excess risk attained by the natural inefficient algorithm scales as $\tilde{O}(d/n)$. It is natural to ask whether this $O(d^{1/2})$ gap can be improved or whether it poses a fundamental barrier. In the slow rate regime, [Barak and Moitra \(2016\)](#) gave a computational lower bound showing that finding efficiently computable classes of tensors for which the Rademacher complexity grows as $o(\sqrt{d^{3/2}/n})$ is at least as hard as refuting random instances of 3-XOR with $o(d^{3/2})$ clauses. In this section we show that this computational hardness is also present in the fast rate regime: Under plausible average-case hardness assumptions, no polynomial time algorithm can obtain a fast rate for square loss scaling as $O(d^{3/2-\varepsilon}/n)$ for any $\varepsilon > 0$.

Our improper learning lower bound applies to *any algorithm* that obtains low excess risk in the sense of [\(15\)](#), and states that under conjectured hardness of a certain distinguishing problem for 3-XOR it is not possible to improve the $O(d^{3/2})$ dependence on dimension.

We reduce from the 3-XOR problem over variables $x \in \{\pm 1\}^d$. A 3-XOR instance consists of a sequence of m clauses of the form

$$x_i \cdot x_j \cdot x_k = z_{ijk},$$

where $z_{ijk} \in \{\pm 1\}$ is a target. We consider two families of instances:

- *Planted.* Fix an arbitrary assignment $a \in \{\pm 1\}^d$. Select m triples (i, j, k) uniformly at random with replacement.⁷ For each such triple (i, j, k) , include a clause

$$x_i \cdot x_j \cdot x_k = z_{ijk} := \begin{cases} a_i \cdot a_j \cdot a_k, & \text{with probability } 1 - \eta \\ -a_i \cdot a_j \cdot a_k, & \text{with probability } \eta. \end{cases}$$

⁷We work in the slightly non-standard with-replacement model to simplify the mapping onto the with-replacement tensor completion model in [Theorem 5](#).

Note, we sample the value $z_{i,j,k}$ for a triple (i, j, k) only once: if the triple is sampled multiple times, the value z_{ijk} will be the same.

- *Random.* Select m triples (i, j, k) uniformly at random with replacement, and take each clause to be $x_i \cdot x_j \cdot x_k = z_{ijk}$, where z_{ijk} is drawn from $\{\pm 1\}$ uniformly at random. Again, we sample the value $z_{i,j,k}$ for a triple (i, j, k) only once.

An algorithm for the distinguishing problem takes m clauses as input and outputs either “Planted” or “Random”. The algorithm is said to succeed if it outputs “Planted” for planted instances and “Random” for random instances with probability at least $1 - o(1)$ over the draw of the instance. Note that the problem becomes easier as η gets smaller, and in particular when $\eta = 0$ the problem can be solved in polynomial time using Gaussian elimination.

Conjecture 1. There is some constant $\eta < 1/4$ such that no algorithm that succeeds for the 3-XOR distinguishing problem with $m = o(d^{3/2})$ runs in polynomial time.

All known polynomial time algorithms for distinguishing require $m = \Omega(d^{3/2})$ clauses, and conjectured hardness of the closely related problem of strong refutation for random 3-XOR for with $o(d^{3/2})$ clauses has been used as a basis to establish hardness of other learning problems (Daniely, 2016; Raghavendra et al., 2017; Kothari et al., 2017; Feldman et al., 2018).

Theorem 7. *Let $\varepsilon > 0$ be fixed. Assuming the 3-XOR distinguishing conjecture, there is no polynomial time algorithm for agnostic tensor completion that guarantees that for any distribution \mathcal{D} , with probability at least $1 - o(1)$,*

$$L_{\mathcal{D}}(\widehat{T}_n) - \inf_{T^*: \text{rank-1}} L_{\mathcal{D}}(T^*) = O\left(\frac{d^{3/2-\varepsilon}}{n}\right). \quad (19)$$

5 Conclusion

Our results demonstrate the power of the sum-of-squares hierarchy for agnostic statistical learning, and show that sum-of-squares algorithms can obtain fast rates for prediction with the square loss. We hope our work will serve as a starting point for applying sum-of-squares to obtain polynomial time algorithms with fast rates in statistical learning for broader classes of models.

A few immediate technical questions emerge. Can the dependence on rank in our results be improved? Can the subgradient results be extended to the general undercomplete or even overcomplete case? Can similar agnostic learning results be obtained with a more practical algorithm that does not rely on solving the full sum-of-squares SDP?

Acknowledgements We thank Sasha Rakhlin and Ankur Moitra for helpful discussions and thank Matthew J. Telgarsky for being a constant source of inspiration.

References

Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory*, pages 417–445, 2016.

Boaz Barak and David Steurer. Sum-of-squares: proofs, beliefs, and algorithms. 2016. URL <https://www.sumofsquares.org/public/lec-definitions-general.html>.

Rémi Bardenet, Odalric-Ambrym Maillard, et al. Concentration inequalities for sampling without replacement. *Bernoulli*, 21(3):1361–1385, 2015.

Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.

Peter L Bartlett, Shahar Mendelson, and Joseph Neeman. ℓ_1 -regularized linear regression: persistence and oracle inequalities. *Probability theory and related fields*, 154(1-2):193–224, 2012.

Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.

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

Amin Coja-Oghlan, Andreas Goerdt, and André Lanka. Strong refutation heuristics for random k-sat. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 310–321. Springer, 2004.

Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 105–117. ACM, 2016.

Kenneth R Davidson and Stanislaw J Szarek. Local operator theory, random matrices and banach spaces. *Handbook of the geometry of Banach spaces*, 1(317-366):131, 2001.

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.

Shmuel Friedland and Lek-Heng Lim. Nuclear norm of higher-order tensors. *Mathematics of Computation*, 87(311):1255–1281, 2018.

Stéphane Gaiffas and Guillaume Lecué. Sharp oracle inequalities for high-dimensional matrix prediction. *IEEE Transactions on Information Theory*, 57(10):6942–6957, 2011.

David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.

Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and Computation*, 100(1):78–150, 1992.

Christopher J Hillar and Lek-Heng Lim. Most tensor problems are np-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.

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.

Prateek Jain and Sewoong Oh. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems*, pages 1431–1439, 2014.

Michael J Kearns, Robert E Schapire, and Linda M Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141, 1994.

Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11(Jul):2057–2078, 2010.

Adam Klivans, Pravesh K Kothari, and Raghu Meka. Efficient algorithms for outlier-robust regression. *Proceedings of The 31st Conference on Learning Theory*, 2018.

Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

Pravesh K Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any csp. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 132–145. ACM, 2017.

Jean B Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

Guillaume Lecué and Shahar Mendelson. Learning subgaussian classes: Upper and minimax bounds. *arXiv preprint arXiv:1305.4825*, 2013.

Guillaume Lecué and Shahar Mendelson. Regularization and the small-ball method ii: complexity dependent error rates. *The Journal of Machine Learning Research*, 18(1):5356–5403, 2017.

Guillaume Lecué and Shahar Mendelson. Regularization and the small-ball method i: sparse recovery. *The Annals of Statistics*, 46(2):611–641, 2018.

Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer-Verlag, New York, 1991.

Tengyuan Liang, Alexander Rakhlin, and Karthik Sridharan. Learning with square loss: Localization through offset rademacher complexity. In *Proceedings of The 28th Conference on Learning Theory*, pages 1260–1285, 2015.

Shahar Mendelson. Learning without concentration. In *Proceedings of The 27th Conference on Learning Theory*, pages 25–39, 2014.

Vitali D Milman and Gideon Schechtman. *Asymptotic theory of finite dimensional normed spaces*. Springer-Verlag New York, Inc., 1986.

Andrea Montanari and Nike Sun. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 2016.

Cun Mu, Bo Huang, John Wright, and Donald Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In *International Conference on Machine Learning*, pages 73–81, 2014.

Sahand Negahban and Martin J Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, pages 1069–1097, 2011.

Sahand Negahban and Martin J Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13(May):1665–1697, 2012.

Sahand Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.

Ryan O’Donnell and Yuan Zhou. Approximability and proof complexity. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1537–1556. Society for Industrial and Applied Mathematics, 2013.

Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.

Aaron Potechin and David Steurer. Exact tensor completion with sum-of-squares. *Proceedings of the 30th Annual Conference on Learning Theory*, 2017.

Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random csps below the spectral threshold. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 121–131. ACM, 2017.

Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12 (Dec):3413–3430, 2011.

R Tyrrell Rockafellar. *Convex Analysis*, volume 28. Princeton University Press, 1970.

Bernardino Romera-Paredes and Massimiliano Pontil. A new convex relaxation for tensor completion. In *Advances in Neural Information Processing Systems*, pages 2967–2975, 2013.

Naum Zuselevich Shor. An approach to obtaining global extrema in polynomial mathematical programming problems. *Cybernetics*, 23(5):695–700, 1987.

Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pages 545–560. Springer, 2005.

Ryota Tomioka and Taiji Suzuki. Convex tensor decomposition via structured schatten norm regularization. In *Advances in neural information processing systems*, pages 1331–1339, 2013.

Ryota Tomioka, Taiji Suzuki, Kohei Hayashi, and Hisashi Kashima. Statistical performance of convex tensor decomposition. In *Advances in Neural Information Processing Systems*, pages 972–980, 2011.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

G Alistair Watson. Characterization of the subdifferential of some matrix norms. *Linear algebra and its applications*, 170:33–45, 1992.

Dong Xia, Ming Yuan, and Cun-Hui Zhang. Statistically optimal and computationally efficient low rank tensor completion from noisy entries. *arXiv preprint arXiv:1711.04934*, 2017.

Ming Yuan and Cun-Hui Zhang. On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068, 2016.

A Preliminaries

A.1 Sum-of-squares proof system

Let $\mathcal{A} = \{f_1 \geq 0, \dots, f_m \geq 0\} \cup \{g_1 = 0, \dots, g_{m'} = 0\}$ be a system of polynomial constraints. A *degree- ℓ sum-of-squares proof* that \mathcal{A} implies a constraint $\{h \geq 0\}$ is a set of polynomials $(q_i)_{i \in [m']}$ and sum-of-squares polynomials $(ps)_{S \subseteq [m]}$ such that

$$h = \sum_{S \subseteq [m]} p_S \prod_{i \in S} f_i + \sum_{i \in [m']} q_i g_i,$$

and where and $\deg(p_S \prod_{i \in S} f_i) \leq \ell$ for all S and $\deg(q_i g_i) \leq \ell$ for all i .⁸ We write $\mathcal{A} \vdash_\ell \{h \geq 0\}$ whenever such a proof exists. Our proofs going forward utilize the well-known duality of SoS proofs and pseudodistributions. See [O'Donnell and Zhou \(2013\)](#) and [Barak and Steurer \(2016\)](#) for further discussion, as well as inference rules for the SoS proof system.

We note the following well-known, but useful lemma:

Lemma 1 (Pseudo-Cauchy Schwarz). Let f and g be polynomials and let $\ell = 2 \max\{\deg f, \deg g\}$. Then for any $\eta > 0$,

$$\vdash_\ell \left\{ fg \leq \frac{\eta}{2} f^2 + \frac{1}{2\eta} g^2 \right\}. \quad (20)$$

As a consequence, if μ is a degree- s pseudodistribution with $s \geq \ell$, then

$$\widetilde{\mathbb{E}}_\mu[fg] \leq \sqrt{\widetilde{\mathbb{E}}_\mu[f^2] \cdot \widetilde{\mathbb{E}}_\mu[g^2]}. \quad (21)$$

A.2 Basic technical results

A.2.1 Sum-of-squares norms

We state a few lemmas capturing useful properties of the sum-of-squares norms.

Proposition 1. The SoS nuclear norm and SoS injective norm are dual: $\|\cdot\|_{\widetilde{\mathbb{E}}_\mu}^* = \|\cdot\|_{\widetilde{\mathbb{E}}_\mu}$ and $\|\cdot\|_{\widetilde{\mathbb{E}}_\mu}^* = \|\cdot\|_{\widetilde{\mathbb{E}}_\mu}$.

Proof of Proposition 1. It is immediate from the norm definitions that

$$\|X\|_{\widetilde{\mathbb{E}}_\mu}^* = \sup_{W \in \mathcal{K}_r} \langle W, X \rangle = \sup_{\substack{\mu \text{ degree-}r \\ \widetilde{\mathbb{E}}_\mu \|x\|_2^2 = \widetilde{\mathbb{E}}_\mu \|y\|_2^2 = \widetilde{\mathbb{E}}_\mu \|z\|_2^2 = 1 \\ W = \widetilde{\mathbb{E}}_\mu[x \otimes y \otimes z]}} \langle W, X \rangle = \|X\|_{\widetilde{\mathbb{E}}_\mu}.$$

The other direction is a consequence of the standard duality theory for finite-dimensional Banach spaces. See, e.g., Theorem 15.4 in [Rockafellar \(1970\)](#). \square

Lemma 2. Let $T = \sum_{i=1}^r \lambda_i \cdot u_i \otimes v_i \otimes w_i$ be an orthogonal rank- r tensor. Then for all $k \geq 4$,

$$\|T\|_{\widetilde{\mathbb{E}}_\mu}^k = \|T\|_{\widetilde{\mathbb{E}}_\mu} = \sum_{i=1}^r |\lambda_i|, \quad \text{and} \quad \|T\|_{\widetilde{\mathbb{E}}_\mu}^k = \|T\|_{\widetilde{\mathbb{E}}_\mu} = \max_{i \leq r} |\lambda_i|. \quad (22)$$

Lemma 2 states that the SoS relaxations of the nuclear norm and injective norm are essentially “integral” for orthogonal tensors. Note, this should not be a huge surprise since it is well-known that polynomial time methods such as power iteration succeed at decomposing orthogonal tensors ([Kolda and Bader, 2009](#)). We should mention it doesn’t seem possible to directly apply such results to give agnostic learning guarantees along the lines of our main theorem. While our benchmark is an orthogonal tensor, the data itself may have no orthogonal structure, and thus there is no clear object to which one might apply such a decomposition.

Proof of Lemma 2. We may assume $\lambda_i \geq 0$ without loss of generality. We first prove equality for the injective norms. Let $i^* = \arg \max_{i \leq r} \lambda_i$. As a starting point, for any k we have $\|T\|_{\widetilde{\mathbb{E}}_\mu}^k \geq \|T\|_{\widetilde{\mathbb{E}}_\mu} \geq \lambda_{i^*}$ by exhibiting $u_{i^*} \otimes v_{i^*} \otimes w_{i^*}$ as a feasible solution to the supremum in $\|T\|_{\widetilde{\mathbb{E}}_\mu} = \sup_{\|x\|_2 = \|y\|_2 = \|z\|_2 = 1} \langle T, x \otimes y \otimes z \rangle$.

⁸We use the convention $\prod_{i \in \emptyset} f_i = 1$, so that if $\vdash_\ell h$ then h itself is a degree- ℓ sum of squares.

For the upper bound, let x, y, z be indeterminates and—exploiting orthogonality—let us change coordinates such that $u_i = v_i = w_i = e_i$. Then we have

$$\langle T, x \otimes y \otimes z \rangle = \sum_{i=1}^r \lambda_i x_i y_i z_i$$

From equation (20), we have

$$\vdash_4 x_i y_i z_i \leq \frac{1}{2} x_i^2 + \frac{1}{2} y_i^2 z_i^2.$$

We also have $\vdash_4 \sum_{i=1}^r y_i^2 z_i^2 \leq (\sum_{i=1}^r y_i^2)(\sum_{i=1}^r z_i^2)$. By the additivity of SoS proofs, and since we have assumed $\lambda_i \geq 0$, this implies

$$\vdash_4 \langle T, x \otimes y \otimes z \rangle \leq \max_{i \leq r} \lambda_i \cdot \left(\frac{1}{2} \|x\|_2^2 + \frac{1}{2} \|y\|_2^2 \|z\|_2^2 \right).$$

Now let $\mathcal{A} = \{\|x\|_2^2 = 1, \|y\|_2^2 = 1, \|z\|_2^2 = 1\}$. We claim $\mathcal{A} \vdash_4 \|y\|_2^2 \|z\|_2^2 = 1$. To see this, write $1 - \|y\|_2^2 \|z\|_2^2 = (1 - \|y\|_2^2)(1 + \|z\|_2^2) + (1 - \|z\|_2^2) + (\|y\|_2^2 - 1)$ and use that $\deg((1 - \|y\|_2^2)(1 + \|z\|_2^2)) = 4$.

Putting everything together, we see that $\mathcal{A} \vdash_4 \langle T, x \otimes y \otimes z \rangle \leq \max_{i \leq r} \lambda_i$. Thus, since the ℓ_2 norm is preserved under change of basis, it follows that if μ is any feasible degree-4 pseudodistribution for the maximization problem (6), we must have $\widetilde{\mathbb{E}}_\mu \langle T, x \otimes y \otimes z \rangle \leq \max_i \lambda_i$, and so $\|T\|_{\widetilde{\text{inj}}_k} \leq \max_i \lambda_i$.

We now establish equality for the nuclear norms. We trivially have $\|T\|_{\widetilde{\text{nuc}}_k} \leq \|T\|_{\text{nuc}} \leq \sum_{i=1}^r \lambda_i$ by exhibiting the decomposition $T = \sum_{i=1}^r \lambda_i \cdot u_i \otimes v_i \otimes w_i$ as a feasible solution to the minimization problem in (4). For the other direction, define $X^* = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$, and observe that the equality we just established for the injective norm implies $\|X^*\|_{\widetilde{\text{inj}}_k} \leq 1$. Thus, using the duality of the SoS nuclear norm and injective norm from Proposition 1, we have

$$\|T\|_{\widetilde{\text{nuc}}_k} = \sup_{X \in (\mathbb{R}^d)^{\otimes 3}: \|X\|_{\widetilde{\text{inj}}_k} \leq 1} \langle X, T \rangle \geq \langle X^*, T \rangle = \sum_{i=1}^r \lambda_i,$$

where the last equality uses that $\{u_i\}$, $\{v_i\}$, and $\{w_i\}$ are all orthogonal. \square

Proposition 2. Let $k \geq 4$. For any degree- k pseudodistribution μ ,

$$\langle T, \widetilde{\mathbb{E}}_\mu [x \otimes y \otimes z] \rangle \leq \|T\|_{\widetilde{\text{inj}}_k} \cdot \sqrt{\widetilde{\mathbb{E}}_\mu \|x\|_2^2 \cdot \widetilde{\mathbb{E}}_\mu \|y\|_2^2 \cdot \widetilde{\mathbb{E}}_\mu \|z\|_2^2}. \quad (23)$$

Furthermore, the following statements hold:

$$\begin{aligned} \{\|z\|_2^2 \leq 1\} \vdash_k \langle T, x \otimes y \otimes z \rangle &\leq \|T\|_{\widetilde{\text{inj}}_k} \cdot \left(\frac{1}{2} \|x\|_2^2 + \frac{1}{2} \|y\|_2^2 \right), \\ \{\|y\|_2^2 \leq 1\} \vdash_k \langle T, x \otimes y \otimes z \rangle &\leq \|T\|_{\widetilde{\text{inj}}_k} \cdot \left(\frac{1}{2} \|x\|_2^2 + \frac{1}{2} \|z\|_2^2 \right), \\ \{\|x\|_2^2 \leq 1\} \vdash_k \langle T, x \otimes y \otimes z \rangle &\leq \|T\|_{\widetilde{\text{inj}}_k} \cdot \left(\frac{1}{2} \|y\|_2^2 + \frac{1}{2} \|z\|_2^2 \right). \end{aligned} \quad (24)$$

Proof of Proposition 2. Equation (23) follows by rescaling a given pseudodistribution μ by using $x' = x / \sqrt{\widetilde{\mathbb{E}}_\mu \|x\|_2^2}$ and so forth, so that the pseudodistribution is feasible for the maximization problem (6).

For (24), let μ be a degree- k pseudodistribution with $\mu \models \{\|z\|_2^2 \leq 1\}$. Then, using (23) and the AM-GM inequality we get $\langle T, \widetilde{\mathbb{E}}_\mu [x \otimes y \otimes z] \rangle \leq \frac{\|T\|_{\widetilde{\text{inj}}_k}}{2} (\widetilde{\mathbb{E}}_\mu \|x\|_2^2 + \widetilde{\mathbb{E}}_\mu \|y\|_2^2)$. Using linearity of the pseudoexpectation operator we have $\widetilde{\mathbb{E}}_\mu \left[\frac{\|T\|_{\widetilde{\text{inj}}_k}}{2} (\|x\|_2^2 + \|y\|_2^2) - \langle T, x \otimes y \otimes z \rangle \right] \geq 0$, and so (24) follows from the duality of pseudoexpectations and sum-of-squares proofs. The remaining statements follow by symmetry. \square

A.2.2 Projections

Here we state some basic results regarding the projection operators defined in [Section 2.2](#).

Proposition 3. Let $x, y, z \in \mathbb{R}^d$ be given. If at least two of the follow conditions hold:

- 1) $x \in \mathbb{U}$,
- 2) $y \in \mathbb{V}$,
- 3) $z \in \mathbb{W}$,

then $\mathcal{Q}_{T^\perp}(x \otimes y \otimes z) = 0$.

Proof of Proposition 3. Suppose that $x \in \mathbb{U}$ and $y \in \mathbb{V}$. Then $\mathcal{P}_{\mathbb{U}^\perp}(x) = 0$, and so $\mathcal{Q}_{T^\perp}^0(x \otimes y \otimes z) = \mathcal{Q}_{T^\perp}^2(x \otimes y \otimes z) = \mathcal{Q}_{T^\perp}^3(x \otimes y \otimes z) = 0$. We also have $\mathcal{P}_{\mathbb{V}^\perp}(y) = 0$, and so $\mathcal{Q}_{T^\perp}^1(x \otimes y \otimes z) = 0$. The remaining cases follow by symmetry. \square

Lemma 3. Let $k \geq 4$. For any tensor $X \in (\mathbb{R}^d)^{\otimes 3}$, and any subspaces $\mathbb{U}, \mathbb{V}, \mathbb{W}$ we have

$$\|(\mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}})X\|_{\tilde{\text{inj}}_k} \leq \|X\|_{\tilde{\text{inj}}_k},$$

and in particular $\|\mathcal{Q}_{T^\parallel}(X)\|_{\tilde{\text{inj}}_k} \leq 4\|X\|_{\tilde{\text{inj}}_k}$ and $\|\mathcal{Q}_{T^\perp}(X)\|_{\tilde{\text{inj}}_k} \leq 4\|X\|_{\tilde{\text{inj}}_k}$.

Proof of Lemma 3. For any degree-4 pseudodistribution μ over indeterminates x, y, z we have

$$\begin{aligned} \tilde{\mathbb{E}}_\mu \langle (\mathcal{P}_{\mathbb{U}} \otimes \mathcal{P}_{\mathbb{V}} \otimes \mathcal{P}_{\mathbb{W}})X, x \otimes y \otimes z \rangle &= \tilde{\mathbb{E}}_\mu \langle X, (\mathcal{P}_{\mathbb{U}}x) \otimes (\mathcal{P}_{\mathbb{V}}y) \otimes (\mathcal{P}_{\mathbb{W}}z) \rangle \\ &\leq \sqrt{\tilde{\mathbb{E}}_\mu \|\mathcal{P}_{\mathbb{U}}x\|_2^2 \cdot \tilde{\mathbb{E}}_\mu \|\mathcal{P}_{\mathbb{V}}y\|_2^2 \cdot \tilde{\mathbb{E}}_\mu \|\mathcal{P}_{\mathbb{W}}z\|_2^2} \\ &\leq \sqrt{\tilde{\mathbb{E}}_\mu \|x\|_2^2 \cdot \tilde{\mathbb{E}}_\mu \|y\|_2^2 \cdot \tilde{\mathbb{E}}_\mu \|z\|_2^2}, \end{aligned}$$

where the first inequality uses [Proposition 2](#) and the second uses that $\|\mathcal{P}_{\mathbb{X}}x\|_2^2 \leq \|x\|_2^2$ for any subspace \mathbb{X} . This establishes the first result.

Now observe from [\(11\)](#) that $\|\mathcal{Q}_{T^\parallel}(T)\|_{\tilde{\text{inj}}_k} \leq \sum_{i=1}^4 \|\mathcal{Q}_{T^\parallel}^i(T)\|_{\tilde{\text{inj}}_k}$ and $\|\mathcal{Q}_{T^\perp}(T)\|_{\tilde{\text{inj}}_k} \leq \sum_{i=1}^4 \|\mathcal{Q}_{T^\perp}^i(T)\|_{\tilde{\text{inj}}_k}$. We thus obtain the second result by applying the first to each of the summands. \square

A.2.3 Flattenings

The multilinear rank of a tensor $T \in (\mathbb{R}^d)^{\otimes 3}$ is the triple (r_1, r_2, r_3) , where

$$r_1(T) = \dim \text{span}\{T_{:,j,k} \mid j, k \in [d]\}, \quad (25)$$

is the dimension of the space spanned by the mode-1 fibers and $r_2(T)$ and $r_3(T)$ are defined likewise for the second and third mode.

We define the i th flattening map $\flat_i : (\mathbb{R}^d)^{\otimes 3} \rightarrow \mathbb{R}^{d \times d^2}$ via

$$\flat_1(T)_{i,(j,k)} = \flat_2(T)_{j,(i,k)} = \flat_3(T)_{k,(i,j)} = T_{i,j,k}. \quad (26)$$

A standard result is that $\text{rank}(\flat_i(T)) = r_i(T)$ ([Friedland and Lim, 2018](#)). We also have the following comparison between the nuclear norm of the tensor and its flattenings.

Lemma 4 ([Friedland and Lim \(2018\)](#), Theorem 9.4). For any tensor $T \in (\mathbb{R}^d)^{\otimes 3}$,

$$\|T\|_{\text{nuc}} \leq \min\left\{\sqrt{\min\{r_2, r_3\}}\|\flat_1(T)\|_{\text{nuc}}, \sqrt{\min\{r_1, r_3\}}\|\flat_2(T)\|_{\text{nuc}}, \sqrt{\min\{r_1, r_2\}}\|\flat_3(T)\|_{\text{nuc}}\right\}.$$

A.2.4 Concentration

Lemma 5 (Talagrand-type concentration for supremum of empirical process). Let \mathcal{F} be a class of functions of the form $f : \mathcal{Z} \rightarrow \mathbb{R}$. Let z_1, \dots, z_n be sampled i.i.d. from a distribution \mathcal{D} over \mathcal{Z} that satisfies $\mathbb{E}[f(z)] = 0$ and has $|f(z)| \leq c$ almost surely. Let $\sigma^2 = \sup_{f \in \mathcal{F}} \mathbb{E} f^2(z)$. Then for any $\delta > 0$, with probability at least $1 - 2\delta$ over the i.i.d. draw of z_1, \dots, z_n ,

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{t=1}^n f(z_t) \right| \leq 4 \mathbb{E}_{z_{1:n}} \mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(z_t) + \sqrt{\frac{2\sigma^2 \log(1/\delta)}{n}} + \frac{2c \log(1/\delta)}{n}.$$

Proof of Lemma 5. Follows from Theorem A.1 of [Bartlett et al. \(2005\)](#) applied to the classes \mathcal{F} and $-\mathcal{F}$ separately, along with the standard in-expectation symmetrization lemma for uniform convergence. \square

B Proofs from Section 2

The main result in this section is to prove [Theorem 4](#), then use this result to prove [Theorem 3](#). Before proceeding to the main proofs we state an intermediate result.

Lemma 6 ([Potechin and Steurer \(2017\)](#)). Let $x, y, z \in \mathbb{R}^d$ be indeterminates. Let $\mathcal{A} = \{\|y\|_2^2 = 1\}$. Then for any $r \in [d]$,

$$\mathcal{A} \vdash_6 \sum_{i=1}^r x_i y_i z_i \leq \frac{1}{2} \|x\|_2^2 + \frac{1}{2} \|z\|_2^2 - \frac{1}{4} \sum_{i=r+1}^d x_i^2 + z_i^2 - \frac{1}{8} \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + z_j^2 + y_j^2 (\|x\|_2^2 + \|z\|_2^2)).$$

Corollary 1. Let $\mathcal{A} = \{\|x\|_2^2 = 1, \|z\|_2^2 = 1\}$. Then for any $r \in [d]$,

$$\mathcal{A} \vdash_6 \sum_{i=1}^r x_i y_i z_i \leq 1 - \frac{1}{4} \sum_{i=r+1}^d x_i^2 + z_i^2 - \frac{1}{8} \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + y_j^2 + z_j^2).$$

Proof of Corollary 1. We will show that $\mathcal{A} \vdash_6 -\sum_{i=1}^d \sum_{j \neq i} y_i^2 y_j^2 \|x\|_2^2 = -\sum_{i=1}^d \sum_{j \neq i} y_i^2 y_j^2$; the term involving $\|z\|_2^2$ follows from the same reasoning. The desired inequality is equivalent to $\sum_{i=1}^d \sum_{j \neq i} y_i^2 y_j^2 (\|x\|_2^2 - 1) = 0$, which is clearly the product of a degree-4 polynomial and the equality constraint $\{\|x\|_2^2 - 1 = 0\}$. \square

Proof of Theorem 4. *Preliminaries.* We first claim that the following equalities hold:

- $\mathcal{Q}_{T^\parallel}^0(X^*) = X^*$.
- $\|X^*\|_{\tilde{\text{inj}}_k} = \|X^*\|_{\tilde{\text{inj}}_{k+2}} = \|X^*\|_{\text{inj}} = 1$.
- $\langle X^*, T \rangle = \|T\|_{\text{nuc}_k} = \|T\|_{\text{nuc}_{k+2}} = \|T\|_{\text{nuc}}$.

Indeed, it is immediate from the definition of X^* that $\mathcal{Q}_{T^\parallel}^0(X^*) = X^*$, and it follows from [Lemma 2](#) that $\|X^*\|_{\tilde{\text{inj}}_k} = 1$ and $\langle X^*, T \rangle = \|T\|_{\text{nuc}_k}$ for all $k \geq 4$.

Bounding dual norm is sufficient. To establish the inequality (13), we reduce to a simpler problem. The claim is as follows: Fix a constant $\alpha > 0$. If for all $X \in (\mathbb{R}^d)^{\otimes 3}$ with $\|X\|_{\tilde{\text{inj}}_k} \leq \alpha$, we have

$$\|X^* + \mathcal{Q}_{T^\perp}(X)\|_{\tilde{\text{inj}}_{k+2}} \leq 1, \tag{27}$$

then (13) holds for all X with $\|X\|_{\widetilde{\text{inj}}_k} \leq \alpha$. To see that this is the case, observe that for any T' and all such X we have

$$\begin{aligned} \langle X^* + \mathcal{Q}_{T^\perp}(X), T' - T \rangle &\leq \|X^* + \mathcal{Q}_{T^\perp}(X)\|_{\widetilde{\text{inj}}_{k+2}} \|T'\|_{\widetilde{\text{nuc}}_{k+2}} - \langle X^* + \mathcal{Q}_{T^\perp}(X), T \rangle \\ &\leq \|T'\|_{\widetilde{\text{nuc}}_{k+2}} - \langle X^* + \mathcal{Q}_{T^\perp}(X), T \rangle \\ &= \|T'\|_{\widetilde{\text{nuc}}_{k+2}} - \|T\|_{\widetilde{\text{nuc}}_{k+2}}, \end{aligned}$$

where the first inequality uses [Proposition 1](#), the second inequality uses (27), and the final equality uses the definition of X^* and that $\langle \mathcal{Q}_{T^\perp}(X), T \rangle = \langle X, \mathcal{Q}_{T^\perp}(T) \rangle = 0$. Rearranging the inequality yields (13).

Bounding the dual norm. The remainder of the proof establishes that (27) holds for $\alpha = 1/64$.

Let $x, y, z \in \mathbb{R}^d$ be indeterminates. We will provide a degree- $(k+2)$ SoS upper bound on the polynomial $\langle X^* + \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle$, which will suffice to establish (27).

Let $\{u_i\}_{i=1}^r, \{v_i\}_{i=1}^r, \{w_i\}_{i=1}^r$ be as in (2). Then, let $\mathbb{U} = \text{span}(\{u_i\}_{i=1}^r)$, $\mathbb{V} = \text{span}(\{v_i\}_{i=1}^r)$, and $\mathbb{W} = \text{span}(\{w_i\}_{i=1}^r)$, so that $\{u_i\}_{i=1}^r$ is a basis for \mathbb{U} , and likewise for the other modes. Let $\{u_i\}_{i=r+1}^d$ be an arbitrary orthonormal basis for \mathbb{U}^\perp and likewise with $\{v_i\}_{i=r+1}^d$ for \mathbb{V}^\perp , and $\{w_i\}_{i=r+1}^d$ for \mathbb{W}^\perp .

We perform a change of basis and let $u_i = v_i = w_i = e_i$, where e_i is the i th standard basis vector. Then with x, y, z expressed in the new basis we can write

$$\langle X^* + \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle = \sum_{i=1}^r x_i y_i z_i + \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle. \quad (28)$$

Let $\mathcal{A} = \{\|x\|_2^2 = 1, \|y\|_2^2 = 1, \|z\|_2^2 = 1\}$. From [Corollary 1](#), we have

$$\mathcal{A} \vdash_6 \sum_{i=1}^r x_i y_i z_i \leq \frac{1}{2} \|x\|_2^2 + \frac{1}{2} \|z\|_2^2 - \frac{1}{4} \sum_{i=r+1}^d x_i^2 + z_i^2 - \frac{1}{8} \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + z_j^2 + y_j^2).$$

We now handle the second term in (28). We will establish that

$$\mathcal{A} \vdash_{k+2} \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq +O(\alpha) \cdot \sum_{i=r+1}^d x_i^2 + z_i^2 + O(\alpha) \cdot \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + z_j^2 + y_j^2). \quad (29)$$

under the assumption that $\|X\|_{\widetilde{\text{inj}}_k} \leq \alpha$. To do this it suffices show that for each i individually,

$$\mathcal{A} \vdash_6 y_i^2 \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq O(\alpha) \cdot y_i^2 \sum_{j \neq i} (x_j^2 + z_j^2 + y_j^2),$$

with an extra additive factor of $O(\alpha) \cdot y_i^2 (x_i^2 + z_i^2)$ when $i > r$.

Let $1 \leq i \leq d$ be fixed and let $x' = x - x_i e_i$, $y' = y - y_i e_i$, and $z' = z - z_i e_i$. We write

$$\langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle = \langle X, \mathcal{Q}_{T^\perp}((x_i e_i + x') \otimes (y_i e_i + y') \otimes (z_i e_i + z')) \rangle. \quad (30)$$

Case: $i \leq r$. Observe that with our change of basis we have $e_i \in \mathbb{U}$ along the first mode, $e_i \in \mathbb{V}$ along the second mode, and $e_i \in \mathbb{W}$ along the third mode. In view of [Proposition 3](#), this means we have

$$\begin{aligned} 0 &= \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes (z_i e_i)) \\ &= \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes z') \\ &= \mathcal{Q}_{T^\perp}(x' \otimes (y_i e_i) \otimes (z_i e_i)) \\ &= \mathcal{Q}_{T^\perp}((x_i e_i) \otimes y' \otimes (z_i e_i)). \end{aligned}$$

Consequently, using multilinearity we can write

$$\begin{aligned} \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle &= \langle \mathcal{Q}_{T^\perp}(X), x' \otimes y' \otimes z' \rangle + \langle \mathcal{Q}_{T^\perp}(X), x' \otimes y' \otimes (z_i e_i) \rangle \\ &\quad + \langle \mathcal{Q}_{T^\perp}(X), (x_i e_i) \otimes y' \otimes z' \rangle + \langle \mathcal{Q}_{T^\perp}(X), x' \otimes (y_i e_i) \otimes z' \rangle. \end{aligned} \quad (31)$$

In what follows we will repeatedly invoke that $\mathcal{A} \vdash_2 \|x_i e_i\|_2^2 \leq \|x\|_2^2 \leq 1$, $\mathcal{A} \vdash_2 \|x'\|_2^2 \leq \|x\|_2^2 \leq 1$, and so forth. We will also use that if $\|X\|_{\text{inj}_k} \leq \alpha$ then—via [Proposition 2](#) and [Lemma 3](#)—for any indeterminates a, b, c ,

$$\begin{aligned} \{\|a\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), a \otimes b \otimes c \rangle &\leq 2\alpha(\|b\|_2^2 + \|c\|_2^2), \\ \{\|b\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), a \otimes b \otimes c \rangle &\leq 2\alpha(\|a\|_2^2 + \|c\|_2^2), \\ \{\|c\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), a \otimes b \otimes c \rangle &\leq 2\alpha(\|b\|_2^2 + \|a\|_2^2). \end{aligned}$$

These inequalities allow us to bound the terms in (31) as follows:

$$\begin{aligned} \{\|y\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x' \otimes y' \otimes z' \rangle &\leq 2\alpha(\|x'\|_2^2 + \|z'\|_2^2), \\ \{\|z\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x' \otimes y' \otimes (z_i e_i) \rangle &\leq 2\alpha(\|x'\|_2^2 + \|y'\|_2^2), \\ \{\|x\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), (x_i e_i) \otimes y' \otimes z' \rangle &\leq 2\alpha(\|y'\|_2^2 + \|z'\|_2^2), \\ \{\|y\|_2^2 \leq 1\} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x' \otimes (y_i e_i) \otimes z' \rangle &\leq 2\alpha(\|x'\|_2^2 + \|z'\|_2^2). \end{aligned}$$

Adding these inequalities, we get

$$\mathcal{A} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha \sum_{j \neq i} (x_j^2 + y_j^2 + z_j^2),$$

and by the multiplication rule for SoS proofs,

$$\mathcal{A} \vdash_{k+2} y_i^2 \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha \sum_{j \neq i} y_i^2 (x_j^2 + y_j^2 + z_j^2).$$

Case: $i > r$. As in the previous case, we split the expression $\langle \mathcal{Q}_{T^\perp}(X), (x_i e_i + x') \otimes (y_i e_i + y') \otimes (z_i e_i + z') \rangle$ in (30) using multilinearity, however we can no longer argue that four of the eight terms vanish. As a starting point, for the four terms that appeared in the $i \leq r$ case we can adopt the same upper bound to get

$$\begin{aligned} \mathcal{A} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle &\leq \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes (z_i e_i)) + \mathcal{Q}_{T^\perp}((x_i e_i) \otimes y' \otimes (z_i e_i)) \\ &\quad + \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes z') + \mathcal{Q}_{T^\perp}(x' \otimes (y_i e_i) \otimes (z_i e_i)) \\ &\quad + 6\alpha \sum_{j \neq i} (x_j^2 + y_j^2 + z_j^2). \end{aligned}$$

We bound the four remaining terms as follows:

$$\begin{aligned} \{\|y\|_2^2 \leq 1\} \vdash_k \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes (z_i e_i)) &\leq 2\alpha(x_i^2 + z_i^2), \\ \{\|y\|_2^2 \leq 1\} \vdash_k \mathcal{Q}_{T^\perp}((x_i e_i) \otimes (y_i e_i) \otimes z') &\leq 2\alpha(x_i^2 + \|z'\|_2^2), \\ \{\|y\|_2^2 \leq 1\} \vdash_k \mathcal{Q}_{T^\perp}(x' \otimes (y_i e_i) \otimes (z_i e_i)) &\leq 2\alpha(\|x'\|_2^2 + z_i^2), \\ \{\|y\|_2^2 \leq 1\} \vdash_k \mathcal{Q}_{T^\perp}((x_i e_i) \otimes y' \otimes (z_i e_i)) &\leq 2\alpha(x_i^2 + z_i^2). \end{aligned}$$

Adding together all of these inequalities, we get

$$\mathcal{A} \vdash_k \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha(x_i^2 + z_i^2) + 8\alpha \sum_{j \neq i} (x_j^2 + y_j^2 + z_j^2),$$

and

$$\mathcal{A} \vdash_{k+2} y_i^2 \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha y_i^2 (x_i^2 + z_i^2) + 8\alpha \sum_{j \neq i} y_i^2 (x_j^2 + y_j^2 + z_j^2).$$

Putting everything together. Taking the inequalities we proved for the individual coordinates i and summing them up, we have

$$\mathcal{A} \vdash_{k+2} \|y\|_2^2 \cdot \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha \sum_{i=r+1}^d y_i^2 (x_i^2 + z_i^2) + 8\alpha \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + y_j^2 + z_j^2).$$

Since $\{\|y\|_2^2 = 1\} \subset \mathcal{A}$, we get

$$\mathcal{A} \vdash_{k+2} \langle \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \leq 6\alpha \sum_{i=r+1}^d (x_i^2 + z_i^2) + 8\alpha \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + y_j^2 + z_j^2).$$

Returning to (28), this inequality plus the earlier bound from [Corollary 1](#) imply

$$\begin{aligned} \mathcal{A} \vdash_{k+2} & \langle X^* + \mathcal{Q}_{T^\perp}(X), x \otimes y \otimes z \rangle \\ & \leq 1 + (6\alpha - 1/4) \sum_{i=r+1}^d (x_i^2 + z_i^2) + (8\alpha - 1/8) \sum_{i=1}^d \sum_{j \neq i} y_i^2 (x_j^2 + z_j^2 + y_j^2) \\ & \leq 1, \quad \text{for } \alpha \leq 1/64. \end{aligned}$$

By the duality of SoS proofs and pseudodistributions, we have $\|X^* + \mathcal{Q}_{T^\perp}(X)\|_{\tilde{\text{inj}}_{k+2}} \leq 1$ as desired. \square

Proof of Theorem 3. We first establish equation (14). We combine the assumption that $\|T'\|_{\widetilde{\text{nuc}}_{k+2}} \leq \|T\|_{\widetilde{\text{nuc}}_k}$ with equation (13) to get

$$\|T\|_{\widetilde{\text{nuc}}_{k+2}} \geq \|T'\|_{\widetilde{\text{nuc}}_{k+2}} \geq \|T\|_{\widetilde{\text{nuc}}_{k+2}} + \langle X^* + \mathcal{Q}_{T^\perp}(X), \Delta \rangle.$$

for all X with $\|X\|_{\tilde{\text{inj}}_k} \leq 1/64$ and X^* as in [Theorem 4](#). Rearranging, this yields

$$\langle X, \mathcal{Q}_{T^\perp}(\Delta) \rangle = \langle \mathcal{Q}_{T^\perp}(X), \Delta \rangle \leq -\langle X^*, \Delta \rangle = -\langle X^*, \mathcal{Q}_{T^\parallel}^0(\Delta) \rangle.$$

We now use that $\|X^*\|_{\tilde{\text{inj}}_k} \leq 1$ (from [Theorem 4](#)) and choose X to be a point obtaining the supremum in $\|\mathcal{Q}_{T^\perp}(\Delta)\|_{\widetilde{\text{nuc}}_k} = \sup_{\|X\|_{\tilde{\text{inj}}_k} \leq 1} \langle X, \mathcal{Q}_{T^\perp}(\Delta) \rangle$, scaled by $1/64$. Then the inequality above implies

$$\langle X, \mathcal{Q}_{T^\perp}(\Delta) \rangle = \frac{1}{64} \cdot \|\mathcal{Q}_{T^\perp}(\Delta)\|_{\widetilde{\text{nuc}}_k} \leq \|\mathcal{Q}_{T^\parallel}^0(\Delta)\|_{\widetilde{\text{nuc}}_k}.$$

We now establish equation (9). Observe that we can write

$$\|\Delta\|_{\widetilde{\text{nuc}}_k} = \|\mathcal{Q}_{T^\parallel}(\Delta) + \mathcal{Q}_{T^\perp}(\Delta)\|_{\widetilde{\text{nuc}}_k} \leq \|\mathcal{Q}_{T^\parallel}(\Delta)\|_{\widetilde{\text{nuc}}_k} + \|\mathcal{Q}_{T^\perp}(\Delta)\|_{\widetilde{\text{nuc}}_k}.$$

Combining this with (14), we get

$$\|\Delta\|_{\widetilde{\text{nuc}}_k} \leq \|\mathcal{Q}_{T^\parallel}(\Delta)\|_{\widetilde{\text{nuc}}_k} + 64 \|\mathcal{Q}_{T^\parallel}^0(\Delta)\|_{\widetilde{\text{nuc}}_k}$$

Using the triangle inequality, we upper bound the first term as

$$\|\mathcal{Q}_{T\|}(\Delta)\|_{\widetilde{\text{nuc}}_k} \leq \sum_{i=1}^4 \|\mathcal{Q}_{T\|}^i(\Delta)\|_{\widetilde{\text{nuc}}_k} \leq \sum_{i=1}^4 \|\mathcal{Q}_{T\|}^i(\Delta)\|_{\text{nuc}}.$$

To proceed, we flatten each tensor in the summation above into a matrix and use [Lemma 4](#) to show that the nuclear norm of the flattening leads to an upper bound. Let $r_1 = \dim \mathbb{U}$, $r_2 = \dim \mathbb{V}$, and $r_3 = \dim \mathbb{W}$. Then the following inequalities hold

$$\begin{aligned} \|\mathcal{Q}_{T\|}^1(\Delta)\|_{\text{nuc}} &\leq \sqrt{r_2} \|\mathbb{b}_3(\mathcal{Q}_{T\|}^1(\Delta))\|_{\text{nuc}} \leq \sqrt{r_2 r_3} \|\mathcal{Q}_{T\|}^1(\Delta)\|_F. \\ \|\mathcal{Q}_{T\|}^2(\Delta)\|_{\text{nuc}} &\leq \sqrt{r_3} \|\mathbb{b}_1(\mathcal{Q}_{T\|}^2(\Delta))\|_{\text{nuc}} \leq \sqrt{r_1 r_3} \|\mathcal{Q}_{T\|}^2(\Delta)\|_F. \\ \|\mathcal{Q}_{T\|}^3(\Delta)\|_{\text{nuc}} &\leq \sqrt{r_1} \|\mathbb{b}_2(\mathcal{Q}_{T\|}^3(\Delta))\|_{\text{nuc}} \leq \sqrt{r_1 r_2} \|\mathcal{Q}_{T\|}^3(\Delta)\|_F. \end{aligned}$$

The first inequality in each line above follows from [Lemma 4](#) and the definitions in (11). The second follows from the fact that $\|A\|_{\text{nuc}} \leq \sqrt{\text{rank}(A)} \|A\|_F$ for any matrix A , along with the fact that $\text{rank}(\mathbb{b}_i(T)) = r_i(T)$ for any tensor, and that flattening does not change the entrywise ℓ_2 norm. We have $\|\mathcal{Q}_{T\|}^0(\Delta)\|_{\text{nuc}} \leq \sqrt{r_2 r_3} \|\mathcal{Q}_{T\|}^0(\Delta)\|_F$ as well by the same argument, though the choice of $r_1/r_2/r_3$ in this case is arbitrary.

To combine all the bounds, we use that $r_1, r_2, r_3 \leq r$ and that orthogonal projection decreases the ℓ_2 norm, which yields

$$\|\Delta\|_{\widetilde{\text{nuc}}_k} \leq 68r \|\Delta\|_F.$$

□

C Proofs from Section 3

This section of the appendix is structured as follows.

First, in [Appendix C.1](#), we provide we provide a generalization bound for general classes of tensors and measurement models, from which all of our main statistical results will follow as special cases. This bound assumes that a restricted eigenvalue-type property holds for the tensor class and measurement model under consideration.

In [Appendix C.2](#) and [Appendix C.3](#) we establish that this restricted eigenvalue property holds for the measurement models in [Section 3](#).

In [Appendix C.4](#) we combine these results to prove the main results of that section.

C.1 Agnostic generalization bounds for tensor classes

In this section we given generalization guarantees for empirical risk minimization in a general learning setup. We receive a set of examples $S := (X_1, Y_1), \dots, (X_n, Y_n)$ i.i.d. from a distribution \mathcal{D} over $(\mathbb{R}^d)^{\otimes 3} \times \mathbb{R}$. We assume that a convex class of tensors $\mathcal{eT} \subseteq (\mathbb{R}^d)^{\otimes 3}$ is given, and that our goal is to achieve excess risk against an unknown benchmark $T^* \in \mathcal{T}$. We analyze the performance of empirical risk minimization over \mathcal{T} :

$$\widehat{T}_n = \arg \min_{T \in \mathcal{T}} \widehat{L}_n(T),$$

where \widehat{L}_n is the empirical square loss. The main result from this section is [Theorem 8](#), which bounds the performance of ERM under various assumptions on the data distribution, the class \mathcal{T} , and the benchmark T^* . To state the result, recall from [Section 3](#) that Σ is the population correlation matrix and \mathcal{X}_n is the empirical design operator.

Theorem 8. *Let a benchmark $T^* \in \mathcal{T}$ be fixed, and let $\xi_t = (Y_t - \langle T^*, X_t \rangle)$. Suppose there exist a pair of dual norms $\|\cdot\|$ and $\|\cdot\|_*$ for which the following conditions hold:*

1. *There are constants $0 \leq c < 2$, $\gamma_n > 0$, and $\delta_0 \geq 0$ such that with probability at least $1 - \delta_0$,*

$$\left\| \Sigma^{1/2} \Delta \right\|_2^F \leq \frac{c}{n} \|\mathcal{X}_n(\Delta)\|_F^2 + \gamma_n \quad \forall \Delta \in \mathcal{T} - T^* \quad (\text{Property 1}).$$

2. *There are constants $M \geq 0$ and $\delta_1 \geq 0$ such that with probability at least $1 - \delta_1$,*

$$\left\| \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\| \leq M \cdot \sqrt{n} \quad (\text{Property 2}).$$

3. *There is a constant $\kappa \geq 0$ such that*

$$\|\Delta\|_*^2 \leq \kappa^2 \cdot \left\| \Sigma^{1/2} \Delta \right\|_F^2 \quad \forall \Delta \in \mathcal{T} - T^* \quad (\text{Property 3}).$$

Then with probability at least $1 - (\delta_0 + \delta_1)$,

$$L_{\mathcal{D}}(\widehat{T}_n) - L_{\mathcal{D}}(T^*) \leq \frac{2\kappa^2 M^2}{c'n} + 2\gamma_n. \quad (32)$$

Proof of Theorem 8. Since \mathcal{T} is convex, and since \widehat{T}_n minimizes the (strongly convex) empirical risk, we have

$$\widehat{L}_n(T^*) - \widehat{L}_n(\widehat{T}_n) \geq \langle \nabla \widehat{L}_n(\widehat{T}_n), T^* - \widehat{T}_n \rangle + \widehat{\mathbb{E}}_n \langle \widehat{T}_n - T^*, X \rangle^2 \geq \widehat{\mathbb{E}}_n \langle \widehat{T}_n - T^*, X \rangle^2.$$

By rearranging and expanding the definition of \widehat{L}_n , this implies

$$\widehat{\mathbb{E}}_n(\langle T^*, X \rangle - Y)^2 - \widehat{\mathbb{E}}_n(\langle \widehat{T}_n, X \rangle - Y)^2 - \widehat{\mathbb{E}}_n \langle \widehat{T}_n - T^*, X \rangle^2 \geq 0.$$

Since the left-hand side is non-negative, we can add it to the population excess risk, which implies

$$\begin{aligned} & \mathbb{E}(\langle \widehat{T}_n, X \rangle - Y)^2 - \mathbb{E}(\langle T^*, X \rangle - Y)^2 \\ & \leq \mathbb{E}(\langle \widehat{T}_n, X \rangle - Y)^2 - \mathbb{E}(\langle T^*, X \rangle - Y)^2 + \widehat{\mathbb{E}}_n(\langle T^*, X \rangle - Y)^2 - \widehat{\mathbb{E}}_n(\langle \widehat{T}_n, X \rangle - Y)^2 - \widehat{\mathbb{E}}_n \langle \widehat{T}_n - T^*, X \rangle^2. \end{aligned}$$

Rearranging, this is equal to

$$= (\mathbb{E} - \widehat{\mathbb{E}}_n) [2 \langle \widehat{T}_n - T^*, X \rangle (\langle T^*, X \rangle - Y)] + \mathbb{E} \langle \widehat{T}_n - T^*, X \rangle^2 - 2 \widehat{\mathbb{E}}_n \langle \widehat{T}_n - T^*, X \rangle^2.$$

Since $\widehat{T}_n - T^*$ is an element of $\mathcal{T} - T^*$, we move to an upper bound by taking a supremum over elements of this set.

$$\leq \sup_{\Delta \in \mathcal{T} - T^*} \{ (\mathbb{E} - \widehat{\mathbb{E}}_n) [2 \langle \Delta, X \rangle (\langle T^*, X \rangle - Y)] + \mathbb{E} \langle \Delta, X \rangle^2 - 2 \widehat{\mathbb{E}}_n \langle \Delta, X \rangle^2 \}.$$

This establishes inequality (18). It remains to use the assumptions in the theorem statement to bound the process. [Property 1](#) states that with probability at least $1 - \delta_0$, $\mathbb{E}\langle \Delta, X \rangle^2 \leq c \cdot \widehat{\mathbb{E}}_n \langle \Delta, X \rangle^2 + \gamma_n$ for all $\Delta \in \mathcal{T} - T^*$. Define $c' = 2 - c > 0$, so that conditioned on this event we have

$$\begin{aligned} & \sup_{\Delta \in \mathcal{T} - T^*} \{(\mathbb{E} - \widehat{\mathbb{E}}_n)[2\langle \Delta, X \rangle(\langle T^*, X \rangle - Y)] + \mathbb{E}\langle \Delta, X \rangle^2 - 2\widehat{\mathbb{E}}_n \langle \Delta, X \rangle^2\} \\ & \leq \sup_{\Delta \in \mathcal{T} - T^*} \{(\mathbb{E} - \widehat{\mathbb{E}}_n)[2\langle \Delta, X \rangle(\langle T^*, X \rangle - Y)] - c' \cdot \widehat{\mathbb{E}}_n \langle \Delta, X \rangle^2\} + \gamma_n. \end{aligned}$$

We expand the first term in the supremum as

$$\begin{aligned} (\mathbb{E} - \widehat{\mathbb{E}}_n)[2\langle \Delta, X \rangle(\langle T^*, X \rangle - Y)] &= -2 \frac{1}{n} \sum_{t=1}^n \xi_t \langle \Delta, X_t \rangle - \mathbb{E}[\xi \langle \Delta, X \rangle] \\ &= -2 \left\langle \frac{1}{n} \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X], \Delta \right\rangle. \end{aligned}$$

It follows from Hölder's inequality that this expression is bounded as

$$2 \left\| \frac{1}{n} \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\| \cdot \|\Delta\|_{\star}.$$

Defining $\psi_n = \left\| \frac{1}{n} \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\|$, the development so far states that with probability at least $1 - \delta_0$,

$$\mathbb{E}(\langle \widehat{T}_n, X \rangle - Y)^2 - \mathbb{E}(\langle T^*, X \rangle - Y)^2 \leq \sup_{\Delta \in \mathcal{T} - T^*} \left\{ 2\psi_n \cdot \|\Delta\|_{\star} - c' \cdot \frac{1}{n} \|\mathcal{X}_n(\Delta)\|_2^2 \right\} + \gamma_n.$$

Using [Property 3](#) we upper bound the leading term by

$$\sup_{\Delta \in \mathcal{T} - T^*} \left\{ 2\kappa\psi_n \cdot \left\| \Sigma^{1/2} \Delta \right\|_F - c' \cdot \frac{1}{n} \|\mathcal{X}_n(\Delta)\|_2^2 \right\},$$

and the event we already conditioned on implies that this is at most

$$\begin{aligned} & \sup_{\Delta \in \mathcal{T} - T^*} \left\{ 4\kappa\psi_n \sqrt{\frac{1}{n} \|\mathcal{X}_n(\Delta)\|_2^2} - c' \cdot \frac{1}{n} \|\mathcal{X}_n(\Delta)\|_2^2 \right\} + 2\kappa\psi_n \sqrt{\gamma_n} \\ & \leq \frac{\kappa^2 \psi_n^2}{c'} + 2\kappa\psi_n \sqrt{\gamma_n}, \end{aligned}$$

where the second inequality follows from AM-GM. Finally, by [Property 2](#) we have that with probability at least $1 - \delta_1$, $\psi_n \leq M/\sqrt{n}$, which leads to the final bound of

$$\frac{\kappa^2 M^2}{c'n} + 2\kappa M \sqrt{\frac{\gamma_n}{n}} + \gamma_n \leq \frac{2\kappa^2 M^2}{c'n} + 2\gamma_n.$$

□

C.2 Restricted eigenvalue for tensor completion

The main result in this section is [Theorem 9](#), which relates the empirical covariance and population covariance for all tensors with bounded entries and bounded SoS nuclear norm under sampling model for tensor completion. This result is then used in the proof of [Theorem 5](#) to establish a restricted eigenvalue guarantee.

Theorem 9. Let $d^{3/2} \leq n \leq d^3$, and let $\varepsilon \in (0, 1/2)$. Suppose observed entries are drawn with replacement. Then for any $\delta > 0$, with probability at least $1 - \delta$, all $T \in (\mathbb{R}^d)^{\otimes 3}$ with $\|T\|_\infty \leq R$ satisfy

$$\|\Sigma^{1/2}T\|_F^2 \leq \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\|T\|_{\text{nuc}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right).$$

The proofs in this section pass back and forth between the with-replacement sampling model for tensor completion used in the main body of the paper and a without-replacement model, in which each entry is only observed a single time for tensor completion. To distinguish between the models, we use $S \sim \mathcal{D}_{\text{wr}}^n$ to refer to the draw of the dataset under the with-replacement sampling model and $S \sim \mathcal{D}_{\text{w/o}}^n$ to refer to the draw under the without-replacement model. $\Omega \subseteq [d]^3$ will denote the support of the set of observed entries.

Before proving [Theorem 9](#), we state and prove a number of auxiliary lemmas, from which the main result will follow.

Proposition 4. Let $N_{j,k} = |\{t \in [n] \mid j_t = j, k_t = k\}|$. Then

$$\mathbb{P}_{S \sim \mathcal{D}_{\text{w/o}}^n} \left(N_{j,k} \geq \frac{2n}{d^2} + \frac{10 \log(1/\delta)}{3} \right) \leq \delta.$$

Proof of Proposition 4. This is an immediate consequence of Bernstein's inequality for without-replacement sampling. See [Bardenet et al. \(2015\)](#), Proposition 1.4. \square

Lemma 7. The Rademacher complexity of $\|\cdot\|_{\widetilde{\text{inj}}_4}$ under without-replacement sampling is bounded as

$$\mathbb{E}_{S \sim \mathcal{D}_{\text{w/o}}^n} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\widetilde{\text{inj}}_4} \leq O\left(\sqrt{\frac{n \log^4 d}{d^{3/2}}} + \sqrt{\log d}\right). \quad (33)$$

Additionally, if $n \leq d^3$, then the Rademacher complexity under with-replacement sampling is bounded as

$$\mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\widetilde{\text{inj}}_4} \leq O\left(\sqrt{\frac{n \log^6 d}{d^{3/2}}} + \sqrt{\log^3 d}\right). \quad (34)$$

Proof of Lemma 7. We first bound the Rademacher complexity in the without-replacement sampling case, then handle the with-replacement case by reduction. This analysis follows [Barak and Moitra \(2016\)](#), except that we handle certain “diagonal” terms that arise in the analysis slightly more carefully so as to get the right scaling for our setup, which differs from theirs in that it does not assume incoherence.

Consider a fixed draw of $\epsilon_{1:n}$ and $X_{1:n}$, and let $Z = \sum_{t=1}^n \epsilon_t X_t$. Letting x, y, z be indeterminates, we will give a degree-four SoS upper bound on the polynomial $\langle Z, x \otimes y \otimes z \rangle$. This will imply that the pseudoexpectation of $\langle Z, x \otimes y \otimes z \rangle$ is bounded for any feasible pseudodistribution in the maximization problem defining $\|\cdot\|_{\widetilde{\text{inj}}_4}$.

To begin, for any fixed constant $\eta > 0$, [Lemma 1](#) implies that

$$\vdash_4 \langle Z, x \otimes y \otimes z \rangle = \sum_{i,j,k} Z_{i,j,k} x_i y_j z_k \leq \frac{1}{2\eta} \sum_{i=1}^d x_i^2 + \frac{\eta}{2} \sum_i \left(\sum_{j,k} Z_{i,j,k} y_j z_k \right)^2,$$

and so

$$\{\|x\|_2^2 = 1\} \vdash_4 \langle Z, x \otimes y \otimes z \rangle \leq \frac{1}{2\eta} + \frac{\eta}{2} \sum_i \left(\sum_{j,k} Z_{i,j,k} y_j z_k \right)^2.$$

Define a matrix $A \in \mathbb{R}^{d^2 \times d^2}$ via $A_{j,k',j',k} = \sum_{i=1}^d Z_{i,j,k} Z_{i,j',k'}$. Define additional matrices $D, B \in \mathbb{R}^{d^2 \times d^2}$ via

$$D_{j,k,j',k'} = \begin{cases} A_{j,k,j,k}, & \text{if } (j,k) = (j',k'), \\ 0, & \text{otherwise,} \end{cases}$$

and $B = A - D$. Then we have

$$\sum_i \left(\sum_{j,k} Z_{i,j,k} y_j z_k \right)^2 = \langle A, (y \otimes z)(y \otimes z)^\top \rangle = \langle B, (y \otimes z)(y \otimes z)^\top \rangle + \langle D, (y \otimes z)(y \otimes z)^\top \rangle.$$

We bound the first term by the operator norm of B via

$$\{\|y\|_2^2 = 1, \|z\|_2^2 = 1\} \vdash_4 \langle B, (y \otimes z)(y \otimes z)^\top \rangle \leq \|B\|_{\text{op}} \|y \otimes z\|_F^2 = \|B\|_{\text{op}} \|y\|_2^2 \|z\|_2^2 \leq \|B\|_{\text{op}}.$$

For the second term, define another matrix $R \in \mathbb{R}^{d \times d}$ via $R_{j,k} = \sum_{i=1}^d Z_{i,j,k}^2$. Then we can write

$$\{\|y\|_2^2 = 1, \|z\|_2^2 = 1\} \vdash_4 \langle D, (y \otimes z)(y \otimes z)^\top \rangle = \sum_{j,k} R_{j,k} y_j^2 z_k^2 \leq \|R\|_\infty \sum_{j,k} y_j^2 z_k^2 = \|R\|_\infty \|y\|_2^2 \|z\|_2^2 \leq \|R\|_\infty.$$

By the duality of sum-of-squares proofs and pseudodistributions, this implies that any degree-four pseudodistribution μ with $\mu \models \{\|x\|_2^2 = 1, \|y\|_2^2 = 1, \|z\|_2^2 = 1\}$ has $\widetilde{\mathbb{E}}\langle Z, x \otimes y \otimes z \rangle \leq \frac{\eta}{2}(\|B\|_{\text{op}} + \|R\|_\infty) + \frac{1}{2\eta}$. Optimizing over η , we conclude that

$$\|Z\|_{\widetilde{\text{inj}}_4} \leq \sqrt{\|B\|_{\text{op}} + \|R\|_\infty}.$$

The bound for the matrix B is taken care of Theorem 4.4 of [Barak and Moitra \(2016\)](#), which implies that $\sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{w/o}}^n} \mathbb{E}_\epsilon \|B\|_{\text{op}}^2} \leq O\left(\frac{n \log^4 d}{d^{3/2}}\right)$. For the matrix R , observe that under the without-replacement sampling model, for any j, k we have

$$R_{j,k} = \sum_{i=1}^d Z_{i,j,k}^2 = \sum_{i=1}^d \mathbb{1}\{(i,j,k) \in \Omega\} = |\{t \in [n] : j_t = j, k_t = k\}|.$$

[Proposition 4](#) thus implies that for any fixed j, k , with probability at least $1 - \delta$,

$$R_{j,k} \leq O\left(n/d^2 + \log(1/\delta)\right),$$

and so, by taking a union bound and integrating out the tail, we have

$$\sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{w/o}}^n} \|R\|_\infty^2} \leq O\left(n/d^2 + \log d\right).$$

Combining the bounds on B and R and using Jensen's inequality yields

$$\mathbb{E}_{S \sim \mathcal{D}_{\text{w/o}}^n} \mathbb{E}_\epsilon \|Z\|_{\widetilde{\text{inj}}_4} \leq \sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{w/o}}^n} \mathbb{E}_\epsilon \|Z\|_{\widetilde{\text{inj}}_4}^2} \leq O\left(\sqrt{\frac{n \log^4 d}{d^{3/2}}} + \sqrt{\log d}\right). \quad (35)$$

This completes the bound for the without-replacement case. For the with-replacement case we reduce to the bound on the second moment above. Condition on the draw of S , and let $T_{i,j,k} = \{t \in [n] : i_t = i, j_t = j, k_t = k\}$. Then we have

$$\mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\tilde{\text{inj}}_4} = \mathbb{E}_\epsilon \left\| \sum_{(i,j,k) \in \Omega} \left(\sum_{t \in T_{i,j,k}} \epsilon_t \right) e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4}.$$

Introduce a new sequence of Rademacher random variables $\sigma \in \{\pm 1\}^{d^3}$. Then the right-hand side above is equal to

$$\mathbb{E}_\epsilon \mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \left(\sum_{t \in T_{i,j,k}} \epsilon_t \right) \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4} \leq \mathbb{E}_\epsilon \max_{i,j,k} \left| \sum_{t \in T_{i,j,k}} \epsilon_t \right| \cdot \mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4},$$

where the inequality follows from the standard Lipschitz contraction lemma for Rademacher complexity (Ledoux and Talagrand, 1991). By the Hoeffding bound, we have that for any fixed index (i,j,k) with probability at least $1 - \delta$,

$$\left| \sum_{t \in T_{i,j,k}} \epsilon_t \right| \leq O \left(\max_{i,j,k} \sqrt{|T_{i,j,k}| \log(1/\delta)} \right).$$

Taking a union bound and integrating out the tail, we have

$$\mathbb{E}_\epsilon \max_{i,j,k} \left| \sum_{t \in T_{i,j,k}} \epsilon_t \right| \leq O \left(\max_{i,j,k} \sqrt{|T_{i,j,k}| \log d} \right).$$

We now move to the final bound by taking the expectation over S . Using Cauchy-Schwarz, the development above implies

$$\mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\tilde{\text{inj}}_4} \leq \sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} |T_{i,j,k}| \log d} \cdot \sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} \mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4}^2}.$$

Bernstein's inequality and the union bound imply that

$$\sqrt{\mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} |T_{i,j,k}| \log d} \leq O \left(\sqrt{\log d (n/d^3 + \log d)} \right) \leq O(\log d),$$

where the second inequality uses that $n \leq d^3$. For the second term, we have

$$\begin{aligned} & \mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} \mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4}^2 \\ &= \sum_{m=1}^n \mathbb{P}_{S \sim \mathcal{D}_{\text{wr}}^n} (|\Omega| = m) \cdot \mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^n} \left[\mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4}^2 \middle| |\Omega| = m \right] \\ &= \sum_{m=1}^n \mathbb{P}_{S \sim \mathcal{D}_{\text{wr}}^n} (|\Omega| = m) \cdot \mathbb{E}_{S \sim \mathcal{D}_{\text{wr}}^m} \mathbb{E}_\sigma \left\| \sum_{(i,j,k) \in \Omega} \sigma_{i,j,k} \cdot e_i \otimes e_j \otimes e_k \right\|_{\tilde{\text{inj}}_4}^2 \\ &\leq \sum_{m=1}^n \mathbb{P}_{S \sim \mathcal{D}_{\text{wr}}^n} (|\Omega| = m) \cdot O \left(\frac{m \log^4 d}{d^{3/2}} + \log d \right) \\ &\leq O \left(\frac{n \log^4 d}{d^{3/2}} + \log d \right), \end{aligned}$$

where the second inequality uses the bound we proved for the without-replacement setting and the second inequality uses $m \leq n$ and that the probabilities sum to one. Combining these two bounds completes the proof. \square

Lemma 8. Define $\mathcal{T} = \{T \in (\mathbb{R}^d)^{\otimes 3} \mid \|T\|_{\text{inj}_4} \leq \tau, \|\Sigma^{1/2}T\|_F \leq \beta, \|T\|_\infty \leq R\}$. Under the with-replacement sampling model, when $d^{3/2} \leq n \leq d^3$, we have that for any $\delta > 0$, with probability at least $1 - \delta$,

$$\left| \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 - \|\Sigma^{1/2}T\|_F^2 \right| \leq O\left(R\tau \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \sqrt{\frac{2R^2\beta^2 \log(1/\delta)}{n}} + \frac{2R^2 \log(1/\delta)}{n} \right) \quad \forall T \in \mathcal{T}.$$

Proof of Lemma 8. To begin, we write

$$\left| \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 - \|\Sigma^{1/2}T\|_F^2 \right| = |\widehat{\mathbb{E}}_n \langle T, X \rangle^2 - \mathbb{E} \langle T, X \rangle^2|.$$

Using this representation, since entries are drawn i.i.d. with replacement, we can apply [Lemma 5](#) with the function class $\mathcal{F} = \{X \mapsto \langle T, X \rangle^2 - \mathbb{E} \langle T, X \rangle^2 \mid T \in \mathcal{T}\}$. In particular, note that for all $T \in \mathcal{T}$ we have $|\langle T, X_t \rangle| = |T_{i_t, j_t, k_t}| \leq R$, and furthermore

$$\sup_{T \in \mathcal{T}} \mathbb{E}(\langle T, X \rangle^2 - \mathbb{E} \langle T, X \rangle^2)^2 \leq \sup_{T \in \mathcal{T}} \mathbb{E} \langle T, X \rangle^4 \leq R^2 \sup_{T \in \mathcal{T}} \mathbb{E} \langle T, X \rangle^2 = R^2 \sup_{T \in \mathcal{T}} \|\Sigma^{1/2}T\|_F^2 \leq R^2 \beta^2.$$

Consequently, [Lemma 5](#) implies that with probability at least $1 - \delta$, for all $T \in \mathcal{T}$,

$$\left| \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 - \|\Sigma^{1/2}T\|_F^2 \right| \leq 4 \mathbb{E}_S \mathbb{E}_\epsilon \sup_{T \in \mathcal{T}} \frac{1}{n} \sum_{t=1}^n \epsilon_t (\langle T, X_t \rangle^2 - \mathbb{E} \langle T, X \rangle^2) + \sqrt{\frac{2R^2\beta^2 \log(1/\delta)}{n}} + \frac{2R^2 \log(1/\delta)}{n}.$$

Using Jensen's inequality and splitting the supremum, we have

$$\mathbb{E}_S \mathbb{E}_\epsilon \sup_{T \in \mathcal{T}} \frac{1}{n} \sum_{t=1}^n \epsilon_t (\langle T, X_t \rangle^2 - \mathbb{E} \langle T, X \rangle^2) \leq 2 \mathbb{E}_S \mathbb{E}_\epsilon \sup_{T \in \mathcal{T}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \langle T, X_t \rangle^2.$$

Using the Lipschitz contraction lemma for Rademacher complexity, we remove the square

$$\mathbb{E}_S \mathbb{E}_\epsilon \sup_{T \in \mathcal{T}} \left[\frac{1}{n} \sum_{t=1}^n \epsilon_t \langle T, X_t \rangle^2 \right] \leq 2R \mathbb{E}_S \mathbb{E}_\epsilon \sup_{T \in \mathcal{T}} \left[\frac{1}{n} \sum_{t=1}^n \epsilon_t \langle T, X_t \rangle \right] \leq \frac{2R\tau}{n} \mathbb{E}_S \mathbb{E}_\epsilon \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t X_t \right\|_{\text{inj}_4}.$$

Finally, using [Lemma 7](#), we have

$$\mathbb{E}_S \mathbb{E}_\epsilon \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t X_t \right\|_{\text{inj}_4} \leq O\left(\sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{\log^{3/2} d}{n} \right) \leq O\left(\sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{\log^{3/2} d}{n} \right).$$

The final bound follows by using $n \geq d^{3/2}$ to simplify this expression. \square

Proof of Theorem 9. Let $\mathcal{T} = \{T \in (\mathbb{R}^d)^{\otimes 3} \mid \|T\|_\infty \leq R\}$. Recall that for tensor completion, the population correlation matrix Σ under with-replacement sampling is equal to $\frac{1}{d^3} I$.

Let $\tau_{\max} = Rd^3$, $\beta_{\max} = Rd^{3/2}$, $\tau_{\min} = Rd^{3/2}/\sqrt{n}$, $\beta_{\min} = R/\sqrt{n}$, and let $N = \lceil \log(\tau_{\max}/\tau_{\min}) \rceil + 1$ and $M = \lceil \log(\beta_{\max}/\beta_{\min}) \rceil + 1$. For each $i \in [N]$ and $j \in [M]$ define $\tau_i = \tau_{\max} e^{1-i}$ and $\beta_j = \beta_{\max} e^{1-j}$. Define

$$\mathcal{T}_{i,j} = \left\{ T \in \mathcal{T} \mid \tau_{i+1} \leq \|T\|_{\text{nuc}_4} \leq \tau_i, \beta_{j+1} \leq \|\Sigma^{1/2}T\|_F \leq \beta_j \right\}.$$

Using [Lemma 8](#) and a union bound, we get that with probability at least $1 - \delta$, for all i, j simultaneously, for all $T \in \mathcal{T}_{i,j}$,

$$\|\Sigma^{1/2}T\|_F^2 \leq \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\tau_i \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \sqrt{\frac{R^2\beta_j^2 \log(MN/\delta)}{n}} + \frac{R^2 \log(MN/\delta)}{n} \right).$$

Now consider a fixed tensor $T \in \mathcal{T}$. There are two cases. First, if $\|T\|_{\text{nuc}_4} \geq \tau_{\min}$ and $\|\Sigma^{1/2}T\|_F \geq \beta_{\min}$, then there must be indices i and j for which $\tau_{i+1} \leq \|T\|_{\text{nuc}} \leq \tau_i$, $\beta_{j+1} \leq \|\Sigma^{1/2}T\|_F \leq \beta_j$. Consequently, the uniform bound above implies

$$\|\Sigma^{1/2}T\|_F^2 \leq \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\|T\|_{\text{nuc}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \|\Sigma^{1/2}T\|_F \sqrt{\frac{R^2 \log(MN/\delta)}{n}} + \frac{R^2 \log(MN/\delta)}{n} \right).$$

On the other hand, if either $\|T\|_{\text{nuc}} \leq \tau_{\min}$ or $\|\Sigma^{1/2}T\|_F \leq \beta_{\min}$, we trivially have

$$\|\Sigma^{1/2}T\|_F^2 \leq \frac{R^2}{n} \leq \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 + \frac{R^2}{n}.$$

Combining these cases, and using the values for N and M , we get that with probability at least $1 - \delta$, for all $T \in \mathcal{T}$,

$$\begin{aligned} \|\Sigma^{1/2}T\|_F^2 &\leq \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\|T\|_{\text{nuc}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \|\Sigma^{1/2}T\|_F \sqrt{\frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{n}} \right) \\ &\quad + O\left(\frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{n} \right). \end{aligned}$$

Using the AM-GM inequality on the second-to-last term and rearranging, we have that for any $\varepsilon > 0$,

$$(1 - \varepsilon) \|\Sigma^{1/2}T\|_F^2 \leq \frac{1}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\|T\|_{\text{nuc}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n} \right).$$

When $\varepsilon \in (0, 1/2)$, this implies

$$\|\Sigma^{1/2}T\|_F^2 \leq \frac{(1 + 2\varepsilon)}{n} \|\mathcal{X}_n(T)\|_2^2 + O\left(R\|T\|_{\text{nuc}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n} \right).$$

□

C.3 Restricted eigenvalue for tensor sensing

In this section we prove the main technical result used to establish restricted eigenvalue guarantees for tensor sensing, which is as follows.

Theorem 10. *Suppose $X \sim \mathcal{N}(0; I)$. There is some universal constant $C > 0$ such that for any $\varepsilon < \sqrt{2/\pi}$, with probability at least $1 - 2e^{-\frac{n}{32}}/(1 - e^{-\frac{n\varepsilon^2}{8}})$,*

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq (\sqrt{2/\pi} - \varepsilon) \|\Delta\|_F - \frac{Cd^{3/4} \log^{1/4} d}{\sqrt{n}} \cdot \|\Delta\|_{\widetilde{\text{inj}}_4} \quad \forall \Delta \in (\mathbb{R}^d)^{\otimes 3}. \quad (36)$$

To prove [Theorem 10](#) we require two key technical lemmas.

Lemma 9. *Let $\|\cdot\|$ be any norm, let $\|\cdot\|_*$ be the dual. Suppose that the rows of \mathcal{X}_n are formed by drawing X_1, \dots, X_n i.i.d. from $\mathcal{N}(0; I_{p \times p})$, and let $\mathbb{E}_{X \sim \mathcal{N}(0; I_{p \times p})} \|X\| \leq \psi$. Then for any $\varepsilon < \sqrt{2/\pi}$, with probability at least $1 - 2e^{-\frac{n}{32}}/(1 - e^{-\frac{n\varepsilon^2}{8}})$,*

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq (\sqrt{2/\pi} - \varepsilon) \|\Delta\|_F - \frac{4\psi}{\sqrt{n}} \|\Delta\|_* \quad \forall \Delta \in \mathbb{R}^p. \quad (37)$$

Lemma 10 (Corollary of [Hopkins et al. \(2015\)](#), Theorem 3.3). *Let $X \in (\mathbb{R}^d)^{\otimes 3}$ have entries drawn i.i.d. from $\mathcal{N}(0; 1)$. Then*

$$\mathbb{E} \|X\|_{\widetilde{\text{inj}}_4} \leq O(d^{3/4} \log^{1/4} d). \quad (38)$$

Proof of Theorem 10. This is an immediate consequence of [Lemma 9](#) and [Lemma 10](#), along with the duality of $\|\cdot\|_{\widetilde{\text{inj}}_4}$ and $\|\cdot\|_{\widetilde{\text{inj}}_4}$. \square

In the remainder of the section we prove [Lemma 9](#). The result follows from fairly standard techniques (e.g. [Wainwright \(2019\)](#)), but we include the proof for completeness. We first restate some basic results on gaussian concentration.

Lemma 11 (Concentration for Lipschitz functions ([Milman and Schechtman, 1986](#))). *Let $Z \in \mathbb{R}^n$ have entries drawn i.i.d. from $\mathcal{N}(0; 1)$, and let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be L -Lipschitz with respect to the ℓ_2 norm. Then for all $t \geq 0$,*

$$\mathbb{P}(|f(Z) - \mathbb{E} f(Z)| \geq t) \leq 2e^{-\frac{t^2}{2L^2}}.$$

Lemma 12 (Gordon's Inequality ([Davidson and Szarek, 2001](#))). *Let $\{Z_{a,b}\}$ and $\{Y_{a,b}\}$ be zero-mean gaussian processes indexed by $A \times B$. Suppose that*

$$\mathbb{E}(Z_{a,b} - Z_{a',b'})^2 \leq \mathbb{E}(Y_{a,b} - Y_{a',b'})^2 \quad \text{for all } (a,b), (a',b') \in A \times B$$

and

$$\mathbb{E}(Z_{a,b} - Z_{a,b'})^2 = \mathbb{E}(Y_{a,b} - Y_{a,b'})^2 \quad \text{for all } a \in A, b, b' \in B.$$

Then

$$\mathbb{E} \sup_{a \in A} \inf_{b \in B} Z_{a,b} \leq \mathbb{E} \sup_{a \in A} \inf_{b \in B} Y_{a,b}.$$

Proof of Lemma 9. *Part 1: Bound at a single scale.*

Define $\mathcal{B}(\tau) = \{\Delta \in \mathbb{R}^p \mid \|\Delta\|_F = 1, \|\Delta\|_* \leq \tau\}$. We will prove that with probability at least $1 - 2e^{-\frac{nt^2}{2}}$,

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq \sqrt{\frac{2}{\pi}} - \frac{\psi\tau}{\sqrt{n}} - t, \quad \forall \Delta \in \mathcal{B}(\tau). \quad (39)$$

We will prove a lower bound on the random variable $\min_{\Delta \in \mathcal{B}(\tau)} \frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2$, which is equivalent to providing an upper bound on the random variable

$$S_n(\tau) = - \inf_{\Delta \in \mathcal{B}(\tau)} \frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 = - \inf_{\Delta \in \mathcal{B}(\tau)} \sup_{\|u\|_2=1} \frac{1}{\sqrt{n}} \langle u, \mathcal{X}_n(\Delta) \rangle = \sup_{\Delta \in \mathcal{B}(\tau)} \inf_{\|u\|_2=1} \frac{1}{\sqrt{n}} \langle u, \mathcal{X}_n(\Delta) \rangle.$$

As a starting point, we show how to upper bound the expectation $\mathbb{E}[S_n(\tau)]$. Define $Z_{\Delta,u} = \frac{1}{\sqrt{n}} \langle u, \mathcal{X}_n(\Delta) \rangle$, so that $S_n(\tau) = \sup_{\Delta \in \mathcal{B}(\tau)} \inf_{\|u\|_2=1} Z_{\Delta,u}$. Note that $Z_{\Delta,u}$ is a gaussian process with variance n^{-1} , since $\|\Delta\|_F = 1$. We now define a new gaussian process that will serve as an upper bound through Gordon's inequality. Let $g \in \mathbb{R}^p$ and $h \in \mathbb{R}^n$ be standard gaussian random variables, and define

$$Y_{\Delta,u} = \frac{1}{\sqrt{n}} \langle g, \Delta \rangle + \frac{1}{\sqrt{n}} \langle h, u \rangle.$$

Note that for any (Δ, u) and (Δ', u') we have

$$\mathbb{E}(Y_{\Delta,u} - Y_{\Delta',u'})^2 = \frac{1}{n} \|\Delta - \Delta'\|_F^2 + \frac{1}{n} \|u - u'\|_2^2.$$

Interpreting Δ and Δ' as vectors in \mathbb{R}^{d^3} , we also have

$$\begin{aligned} \mathbb{E}(Z_{\Delta,u} - Z_{\Delta',u'})^2 &= \frac{1}{n} \|u\Delta^\top - u'\Delta'^\top\|_F^2 \\ &= \frac{1}{n} \|\Delta - \Delta'\|_F^2 + \frac{1}{n} \|u - u'\|_2^2 + \frac{1}{n} \langle (u - u')\Delta^\top, u'(\Delta - \Delta')^\top \rangle \\ &= \frac{1}{n} \|\Delta - \Delta'\|_F^2 + \frac{1}{n} \|u - u'\|_2^2 + \frac{1}{n} (\langle u, u' \rangle - 1)(1 - \langle \Delta, \Delta' \rangle) \\ &\leq \frac{1}{n} \|\Delta - \Delta'\|_F^2 + \frac{1}{n} \|u - u'\|_2^2, \end{aligned}$$

where we have used that $\|u\|_2 = \|u'\|_2 = \|\Delta\|_F = \|\Delta'\|_F = 1$. It is also easily seen from the representation above that for any triple (Δ, u, u') we have equality: $\mathbb{E}(Z_{\Delta,u} - Z_{\Delta,u'})^2 = \mathbb{E}(Y_{\Delta,u} - Y_{\Delta,u'})^2$. This means that the preconditions of [Lemma 12](#) are satisfied, and so

$$\begin{aligned} \mathbb{E}[S_n(\tau)] &\leq \mathbb{E} \sup_{\Delta \in \mathcal{B}(\tau)} \inf_{\|u\|_2=1} Y_{\Delta,u} \\ &= \frac{1}{\sqrt{n}} \mathbb{E}_g \sup_{\Delta \in \mathcal{B}(\tau)} \langle g, \Delta \rangle + \mathbb{E}_h \frac{1}{\sqrt{n}} \min_{\|u\|_2=1} \langle h, u \rangle \\ &= \frac{\tau}{\sqrt{n}} \cdot \mathbb{E}_g \|g\| - \frac{1}{\sqrt{n}} \mathbb{E}_h \|h\|_2 \\ &\leq \frac{\psi\tau}{\sqrt{n}} - \sqrt{\frac{2}{\pi}}. \end{aligned}$$

The provides the desired upper bound in expectation. To establish the high probability result we appeal to gaussian concentration for Lipschitz functions. Let $\mathbf{X} \in \mathbb{R}^{n \times d^3}$ denote the sequence of measurements X_1, \dots, X_n , interpreted as a matrix with vectorized measurements as rows. Define $f : \mathbb{R}^{n \times d^3} \rightarrow \mathbb{R}$ via $f(\mathbf{X}) = \min_{\Delta \in \mathcal{B}(\tau)} \frac{1}{\sqrt{n}} \|\mathbf{X}\Delta\|_2$. Observe that we have

$$|f(\mathbf{X}) - f(\mathbf{X}')| \leq \sup_{\Delta \in \mathcal{B}(\tau)} \frac{1}{\sqrt{n}} \|(\mathbf{X} - \mathbf{X}')\Delta\|_2 \leq \frac{1}{\sqrt{n}} \|\mathbf{X} - \mathbf{X}'\|_{\text{op}} \leq \frac{1}{\sqrt{n}} \|\mathbf{X} - \mathbf{X}'\|_F,$$

and so f is $n^{-\frac{1}{2}}$ -Lipschitz with respect to ℓ_2 . [Lemma 11](#) therefore implies that

$$\mathbb{P}(|S_n(\tau) - \mathbb{E} S_n(\tau)| \geq t) \leq 2e^{-\frac{nt^2}{2}}.$$

In other words, [\(39\)](#) holds.

Part 2: Bound at all scales. We will show that for any $\varepsilon < \sqrt{2/\pi}$, with probability at least $1 - 2e^{-\frac{n}{32}}/(1 - e^{-\frac{n\varepsilon^2}{8}})$,

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \geq \sqrt{\frac{2}{\pi}} - \varepsilon - \frac{4\psi}{\sqrt{n}}\|\Delta\|_*, \quad \forall \Delta : \|\Delta\|_F = 1. \quad (40)$$

Define $\mathcal{B}(\tau_\ell, \tau_u) = \{\Delta \in \mathbb{R}^p \mid \|\Delta\|_F = 1, \tau_\ell \leq \frac{\psi}{\sqrt{n}}\|\Delta\| \leq \tau_u\}$. Set $\mu = \varepsilon/2$ and consider the classes $\mathcal{B}(0, \mu)$ and $\mathcal{B}(2^{i-1}\mu, 2^i\mu)$ for $i \in \mathbb{N}$. Note that if equation [\(40\)](#) fails to hold and $\Delta \in \mathcal{B}(0, \mu)$ then

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \leq \sqrt{\frac{2}{\pi}} - \varepsilon - \frac{4\psi}{\sqrt{n}}\|\Delta\|_* \leq \sqrt{\frac{2}{\pi}} - \varepsilon = \sqrt{\frac{2}{\pi}} - 2\mu.$$

Furthermore, if equation [\(40\)](#) fails to hold and $\Delta \in \mathcal{B}(2^{i-1}\mu, 2^i\mu)$, then

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \leq \sqrt{\frac{2}{\pi}} - \varepsilon - \frac{4\psi}{\sqrt{n}}\|\Delta\|_* \leq \sqrt{\frac{2}{\pi}} - 2 \cdot 2^i\mu.$$

Our development in part 1 of the proof implies that for any fixed $\tau_\ell \leq \tau_u$, with probability at least $1 - 2e^{-\frac{n}{32}}e^{-\frac{\tau_u^2}{2}}$,

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \geq \sqrt{\frac{2}{\pi}} - 2\tau_u, \quad \forall \Delta \in \mathcal{B}(\tau_\ell, \tau_u). \quad (41)$$

Thus, by a union bound, we get that with probability at least $1 - 2e^{-\frac{n}{32}}\sum_{i=0}^{\infty} e^{-\frac{2^{2i}n\varepsilon^2}{8}}$, or conservatively at least $1 - 2e^{-\frac{n}{32}}/(1 - e^{-\frac{n\varepsilon^2}{8}})$,

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \geq \sqrt{\frac{2}{\pi}} - 2 \cdot 2^i\mu, \quad \forall \Delta \in \mathcal{B}(2^{i-1}\mu, 2^i\mu), \quad \forall i \in \mathbb{N},$$

and

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \geq \sqrt{\frac{2}{\pi}} - 2\mu, \quad \forall \Delta \in \mathcal{B}(0, \mu),$$

or in other words, equation [\(40\)](#) holds.

We extend the guarantee in equation [\(40\)](#) to arbitrary $\Delta \in \mathbb{R}^p$ by rescaling so that $\|\Delta\|_F = 1$ and then exploiting homogeneity to get

$$\frac{1}{\sqrt{n}}\|\mathcal{X}_n(\Delta)\|_2 \geq \left(\sqrt{\frac{2}{\pi}} - \varepsilon\right)\|\Delta\|_F - \frac{4\psi}{\sqrt{n}}\|\Delta\|_*, \quad \forall \Delta \in \mathbb{R}^p.$$

□

C.4 Proofs of main results

Proof of Theorem 5. We prove the theorem by appealing to the generic result of [Theorem 8](#) using norms $\|\cdot\| = \|\cdot\|_{\widetilde{\text{inj}}_4}$ and $\|\cdot\|_* = \|\cdot\|_{\widetilde{\text{nuc}}_4}$. Let T^* be an arbitrary rank- r orthogonal tensor with $\|T^*\|_{\widetilde{\text{nuc}}_6} = \tau$ and $\|T^*\|_\infty \leq R$.

First, observe that all elements $T \in \mathcal{T}$ have $\|T\|_{\widetilde{\text{nuc}}_6} \leq \|T^*\|_{\widetilde{\text{nuc}}_6}$. Consequently, [Theorem 3](#) implies that all $\Delta \in \mathcal{T} - T^*$ satisfy

$$\|\Delta\|_{\widetilde{\text{nuc}}_4} \leq 68r \cdot \|\Delta\|_F, \quad (42)$$

which establishes [Property 3](#) with $\kappa = O(r \cdot d^{3/2})$.

To establish [Property 2](#) we appeal to [Theorem 9](#), which implies that for any $\varepsilon < 1/2$, with probability at least $1 - \delta$, all $\Delta \in \mathcal{T} - T^*$ satisfy

$$\|\Sigma^{1/2}\Delta\|_F^2 \leq \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R\|\Delta\|_{\widetilde{\text{nuc}}_4} \sqrt{\frac{\log^6 d}{nd^{3/2}}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right).$$

Using equation (42) this is upper bounded by

$$\begin{aligned} \|\Sigma^{1/2}\Delta\|_F^2 &\leq \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R\|\Delta\|_F \sqrt{\frac{r^2 \log^6 d}{nd^{3/2}}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right) \\ &= \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R\|\Sigma^{1/2}\Delta\|_F \sqrt{\frac{r^2 d^{3/2} \log^6 d}{n}} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right). \end{aligned}$$

Using the AM-GM inequality, this is further upper bounded by

$$\|\Sigma^{1/2}\Delta\|_F^2 \leq \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + \varepsilon \|\Sigma^{1/2}\Delta\|_F^2 + O\left(R^2 \frac{r^2 d^{3/2} \log^6 d}{n} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right).$$

Rearranging, this is equivalent to

$$(1 - \varepsilon) \|\Sigma^{1/2}\Delta\|_F^2 \leq \frac{(1+2\varepsilon)}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R^2 \frac{r^2 d^{3/2} \log^6 d}{n} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right),$$

and since $\varepsilon < 1/2$ this implies

$$\|\Sigma^{1/2}\Delta\|_F^2 \leq \frac{(1+2\varepsilon)^2}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R^2 \frac{r^2 d^{3/2} \log^6 d}{n} + \frac{R^2 \log(\log^2(d^{3/2}\sqrt{n})/\delta)}{\varepsilon n}\right).$$

Making the somewhat arbitrary choice of $\varepsilon = 1/100$, and simplifying the right-hand-side, we get

$$\|\Sigma^{1/2}\Delta\|_F^2 \leq \frac{1.1}{n} \|\mathcal{X}_n(\Delta)\|_2^2 + O\left(R^2 \frac{r^2 d^{3/2} \log^6 d + \log(1/\delta)}{n}\right).$$

So we can take $c = 1.1$ and $\gamma_n = O\left(R^2 \frac{r^2 d^{3/2} \log^6 d + \log(1/\delta)}{n}\right)$.

Finally, we establish [Property 1](#). [Lemma 5](#) establishes that with probability at least $1 - \delta$,

$$\left\| \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\|_{\widetilde{\text{inj}}_4} \leq 4 \mathbb{E}_S \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t (\langle T^*, X_t \rangle - Y_t) X_t \right\|_{\widetilde{\text{inj}}_4} + \sqrt{2\sigma^2 n \log(1/\delta)} + 4R \log(1/\delta),$$

where $\sigma^2 = \sup_{\|T\|_{\widetilde{\text{nuc}}_4} \leq 1} \mathbb{E}[\langle T, X \rangle^2 (\langle T^*, X_t \rangle - Y_t)^2] \leq \sup_{\|T\|_{\widetilde{\text{nuc}}_4} \leq 1} \frac{4R^2}{d^3} \|T\|_F^2 \leq \frac{4R^2}{d^3}$. Using [Lemma 7](#) with the standard in-expectation Lipschitz contraction lemma (e.g., [Ledoux and Talagrand \(1991\)](#)), we have

$$4\mathbb{E}_S \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t (\langle T^*, X_t \rangle - Y_t) X_t \right\|_{\widetilde{\text{inj}}_4} \leq 8R \mathbb{E}_S \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\widetilde{\text{inj}}_4} \leq O\left(R \sqrt{\frac{n \log^6 d}{d^{3/2}}} + R \sqrt{\log^3 d} \right).$$

Since $n = \Omega(d^{3/2})$, these bounds together imply that we can take $M = O\left(R \sqrt{\frac{\log^6 d + \log(1/\delta)}{d^{3/2}}} \right)$.

[Theorem 8](#) now implies the claimed result. \square

Proof of Theorem 6. As in the tensor completion case, we prove the theorem by appealing to [Theorem 8](#) using norms $\|\cdot\| = \|\cdot\|_{\widetilde{\text{inj}}_4}$ and $\|\cdot\|_* = \|\cdot\|_{\widetilde{\text{nuc}}_4}$.

Let T^* be an arbitrary orthogonal tensor with $\|T^*\|_{\widetilde{\text{nuc}}_6} = \tau$, $R(T^*) = R$, and $r(T^*) = r$. [Theorem 3](#) implies that all $\Delta \in \mathcal{T} - T^*$ satisfy

$$\|\Delta\|_{\widetilde{\text{nuc}}_4} \leq 68r \cdot \|\Delta\|_F \quad (43)$$

and thus, since $\Sigma = I$, we may take $\kappa = 68r$ to establish [Property 3](#).

To establish [Property 2](#) we appeal to [Theorem 10](#). This implies that for any $\varepsilon < \sqrt{2/\pi}$, with probability at least $1 - 2e^{-\frac{n}{32}}/(1 - e^{-\frac{n\varepsilon^2}{8}})$,

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq (\sqrt{2/\pi} - \varepsilon) \|\Delta\|_F - \frac{Cd^{3/4} \log^{1/4} d}{\sqrt{n}} \cdot \|\Delta\|_{\widetilde{\text{nuc}}_k},$$

where $C > 0$ is some absolute constant. Using equation (43), we have that for all $\Delta \in \mathcal{T} - T^*$, conditioned on the event above,

$$\frac{1}{\sqrt{n}} \|\mathcal{X}_n(\Delta)\|_2 \geq (\sqrt{2/\pi} - \varepsilon) \|\Delta\|_F - \frac{C' r d^{3/4} \log^{1/4} d}{\sqrt{n}} \cdot \|\Delta\|_F.$$

This means that when $n = \Omega(r^2(T^*)d^{3/2} \log^{1/2} d/\varepsilon)$, we have

$$\|\Delta\|_F^2 \leq \frac{1}{(\sqrt{2/\pi} - 2\varepsilon)^2} \frac{1}{n} \|\mathcal{X}_n(\Delta)\|_2^2.$$

It suffices to set $\varepsilon = 1/40$ to get

$$\|\Delta\|_F^2 \leq \frac{1.8}{n} \|\mathcal{X}_n(\Delta)\|_2^2.$$

So, simplifying, [Property 2](#) is satisfied with $c = 1.8$ and $\gamma_n = 0$ with probability at least $1 - e^{-\frac{n}{64}}$ when $n = \Omega(r^2(T^*)d^{3/2} \log^{1/2} d)$.

To establish [Property 1](#) we use [Lemma 10](#), but some care needs to be taken to establish that this applies with high probability. Pick a constant $\tau \geq 0$, and observe via [Lemma 11](#) that

$$\mathbb{P}\left(\|X\|_F \geq d^{3/2} + \tau\right) \leq 2e^{-\frac{\tau^2}{2}}.$$

Define a truncated sequence $X'_t = X_t \mathbb{1}\{\|X_t\|_F \leq d^{3/2} + \tau\}$, and set $\delta_0 = 2ne^{-\frac{\tau^2}{2}}$. Observe that with probability at least $1 - \delta_0$, $X_t = X'_t$ for all t , and so

$$\left\| \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\|_{\tilde{\text{inj}}_4} \leq \left\| \sum_{t=1}^n \xi_t X'_t - \mathbb{E}[\xi X'] \right\|_{\tilde{\text{inj}}_4}.$$

We apply [Lemma 5](#) to the truncated complexity, which implies that with probability at least $1 - (\delta + \delta_0)$,

$$\begin{aligned} & \left\| \sum_{t=1}^n \xi_t X'_t - \mathbb{E}[\xi X'] \right\|_{\tilde{\text{inj}}_4} \\ & \leq 4 \mathbb{E}_{X_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t (\langle T^*, X_t \rangle - Y_t) X'_t \right\|_{\tilde{\text{inj}}_4} + O\left(\sqrt{\sigma^2 n \log(1/\delta)} + R(d^{3/2} + \tau) \cdot \log(1/\delta)\right), \end{aligned}$$

where $\sigma^2 = \sup_{\|T\|_{\text{inj}_4} \leq 1} \mathbb{E}[\langle T, X' \rangle^2 (\langle T^*, X_t \rangle - Y_t)^2] \leq \sup_{\|T\|_{\text{inj}_4} \leq 1} \|T\|_F^2 \cdot 4R^2 \leq 4R^2$. We now apply Lipschitz contraction to bound the in-expectation Rademacher complexity as

$$4 \mathbb{E}_{X_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t (\langle T^*, X_t \rangle - Y_t) X'_t \right\|_{\tilde{\text{inj}}_4} \leq 8R \mathbb{E}_{X'_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X'_t \right\|_{\tilde{\text{inj}}_4}.$$

Integrating out the tail, we can bound the error due to truncation as

$$\mathbb{E}_{X_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X'_t \right\|_{\tilde{\text{inj}}_4} \leq \mathbb{E}_{X_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\tilde{\text{inj}}_4} + n \cdot \int_{d^{3/2} + \tau}^\infty \mathbb{P}(\|X\|_F > t) dt$$

We have

$$n \cdot \int_{d^{3/2} + \tau}^\infty \mathbb{P}(\|X\|_F > t) dt = n \cdot \int_{d^{3/2} + \tau}^\infty e^{-\frac{t^2}{2}} dt \leq n \cdot \exp(-C(d^3 + \tau^2)),$$

for some absolute constant C . We pick $\tau = (\sqrt{(d^{3/2}n + \log(1/\delta))/C})$, so that the quantity above is $o(1)$ and $\delta_0 \leq \delta$. Lastly, since $X_{1:n}$ are gaussian, [Lemma 10](#) implies

$$\mathbb{E}_{X_{1:n}} \mathbb{E}_\epsilon \left\| \sum_{t=1}^n \epsilon_t X_t \right\|_{\tilde{\text{inj}}_4} = \sqrt{n} \cdot \mathbb{E}_X \|X\|_{\tilde{\text{inj}}_4} \leq O(d^{3/4} \log^{1/4} d \sqrt{n}).$$

When $n \geq d^{3/2}$, we have $d^{3/2} \leq d^{3/4} \sqrt{n}$, and so we conclude that with probability at least $1 - \delta$,

$$\left\| \sum_{t=1}^n \xi_t X_t - \mathbb{E}[\xi X] \right\|_{\tilde{\text{inj}}_4} \leq O(R d^{3/4} \log^{3/2} (d/\delta) \sqrt{n}),$$

so $M = O(R d^{3/4} \log^{3/2} (d/\delta))$.

□

D Proofs from [Section 4](#)

Proof of [Theorem 7](#). Let $\varepsilon > 0$ be fixed and let $m = d^{3/2 - \varepsilon/2}$.

We will generate an instance of tensor completion from the 3-XOR instance by treating the indices (i, j, k) in each clause as an observed entry. Precisely, for each clause, we set the corresponding $X = e_i \otimes e_j \otimes e_k$,

and the corresponding $Y = z_{ijk}$. The induced risk for each tensor T in this setting is simply $L_{\mathcal{D}_Z}(T) := \frac{1}{d^3} \sum_{i,j,k} (T_{i,j,k} - Z_{i,j,k})^2$, where the tensor $Z \in (\{\pm 1\}^d)^{\otimes 3}$ is defined as follows: in the random case, the entries of Z are selected uniformly at random from $\{\pm 1\}$; in the planted case with planted assignment $a \in \{\pm 1\}^d$, we start off with the tensor $a^{\otimes 3}$, then get Z by flipping each coordinate independently with probability η .

If we set $n = \Omega(m)$, then we are guaranteed to receive $\tilde{\Omega}(m)$ unique clauses under with-replacement sampling, since there are at most logarithmically many repeats with inverse polynomial probability for $m = o(d^{3/2})$. With this observation, it is clear that any algorithm that takes as input the examples S and outputs “Planted” with probability at least $1 - o(1)$ when Z comes from the planted distribution and “Random” with probability at least $1 - o(1)$ when Z is from the random distribution is a successful distinguisher for the planted 3-XOR problem with $\tilde{\Theta}(m)$ clauses. Going forward we assume $n = \Omega(d)$, since this is the interesting regime for distinguishing and refutation.

Now, suppose we have an algorithm that enjoys the excess risk bound (19) for any n , and let \widehat{T}_S denote its output given input dataset S . Since $Z \in (\{\pm 1\}^d)^{\otimes 3}$ we can take $\|\widehat{T}_S\|_\infty \leq 1$ without loss of generality. We will turn \widehat{T}_S into a successful distinguishing algorithm for the planted 3-XOR problem as follows: Define $\gamma = 1 - 4\eta$, and note the assumption that $\eta < 1/4$ implies that $\gamma > 0$. Split the sample set S into halves S_1 and S_2 , and let \widehat{L}_{S_1} and \widehat{L}_{S_2} denote the empirical risk on the respective halves. Let \widehat{T}_{S_1} be the output of the assumed algorithm on S_1 . If $\widehat{L}_{S_2}(\widehat{T}_{S_1}) \geq 1 - \gamma^4/2$ return “Random”, else return “Planted”.

We will show that the algorithm succeeds in both the planted and random case by analyzing the value of $\widehat{L}_{S_2}(\widehat{T}_{S_1})$ in each case.

Random case. Let S'_2 be the result of removing all entries that appear in S_1 from S_2 . The number of repeats is at most $O(n^2/d^3 + \log(1/\delta))$ with probability at least $1 - \delta$. It follows that with probability at least, say, $1 - O(d^{-1})$,

$$\frac{|S'_2|}{|S_2|} \geq 1 - o(1), \quad \text{and} \quad \widehat{L}_{S_2}(\widehat{T}_{S_1}) \geq \widehat{L}_{S'_2}(\widehat{T}_{S_1}) - o(1).$$

Observe that for every remaining example $X = e_i \otimes e_j \otimes e_k$ in S'_2 , the value of \widehat{T}_{S_1} is statistically independent of the value of $Z_{i,j,k}$. Abbreviating \widehat{T}_{S_1} to \widehat{T} , this means that we have

$$\mathbb{E}_Z[\widehat{L}_{S'_2}(\widehat{T}_{S_1})] = \frac{1}{|S'_2|} \sum_{X=e_i \otimes e_j \otimes e_k \in S'_2} \mathbb{E}_{Z_{i,j,k} \in \{\pm 1\}} (\widehat{T}_{i,j,k} - Z_{i,j,k})^2 = \frac{1}{|S'_2|} \sum_{X=e_i \otimes e_j \otimes e_k \in S'_2} (\widehat{T}_{i,j,k})^2 + 1 \geq 1.$$

Condition on S'_2 . Since \widehat{T}_{S_1} and Z have bounded entries, it follows from Hoeffding’s inequality that with probability at least $1 - \delta$ over the choice of Z ,

$$\widehat{L}_{S'_2}(\widehat{T}_{S_1}) \geq \mathbb{E}_Z[\widehat{L}_{S'_2}(\widehat{T}_{S_1})] - c \sqrt{\frac{\log(1/\delta)}{|S'_2|}},$$

for absolute constant $c > 0$. Since $|S'_2| = \tilde{\Omega}(m)$, we can use the AM-GM inequality to conclude with probability at least $1 - O(d^{-1})$,

$$\widehat{L}_{S_2}(\widehat{T}_{S_1}) \geq \widehat{L}_{S'_2}(\widehat{T}_{S_1}) - o(1).$$

It follows by union bound that

$$\widehat{L}_{S_2}(\widehat{T}_{S_1}) \geq 1 - o(1) \geq 1 - \gamma^4/100 - o(1).$$

This proves that our strategy will indeed return “Random” for random instances.

Planted case. For any Z , the assumed excess risk bound implies that for any rank-1 tensor T^* with bounded entries, we have

$$L_{\mathcal{D}_Z}(\widehat{T}_{S_1}) - L_{\mathcal{D}_Z}(T^*) \leq o(1),$$

when $n = \omega(d^{3/2-\varepsilon})$. We choose $T^* = a^{\otimes 3}$, where $a \in \{\pm 1\}^d$ is the planted assignment. Taking expectation over the flips in Z , we have

$$\mathbb{E}_Z[L_{\mathcal{D}_Z}(a^{\otimes 3})] = \mathbb{E}_Z\left[\frac{4}{d^3} \sum_{i,j,k} \mathbb{1}\{Z_{ijk} \text{ flipped}\}\right] = 4\eta.$$

Applying Bernstein's inequality, we have that with probability at least $1 - O(d^{-1})$ over the choice of Z ,

$$L_{\mathcal{D}_Z}(a^{\otimes 3}) \leq 4(1 + \gamma)\eta + o(1) = 1 - \gamma^2 + o(1).$$

Finally, we use Bernstein once more to show that the empirical loss for S_2 converges to the population loss, leveraging that S_2 is an independent hold-out set for \widehat{T}_{S_1} . We have that for any choice of S_1 , with probability at least $1 - \delta$ over the draw of S_2 ,

$$\widehat{L}_{S_2}(\widehat{T}_{S_1}) \leq L_{\mathcal{D}_Z}(\widehat{T}_{S_1}) + c_1 \sqrt{\frac{L_{\mathcal{D}_Z}(\widehat{T}_{S_1}) \log(1/\delta)}{n}} + c_2 \frac{\log(1/\delta)}{n},$$

where c_1 and c_2 are absolute constants. Applying AM-GM, this implies that $\widehat{L}_{S_2}(\widehat{T}_{S_1}) \leq (1 + \gamma^2)L_{\mathcal{D}_Z}(\widehat{T}_{S_1}) + o(1)$ with probability at least $1 - O(d^{-1})$, and so by combining this with the excess risk bound and the bound on $L_{\mathcal{D}_Z}(a^{\otimes 3})$, we have

$$\widehat{L}_{S_2}(\widehat{T}_{S_1}) \leq 1 - \gamma^4 + o(1).$$

□