

A Non-generative Framework and Convex Relaxations for Unsupervised Learning

Elad Hazan Tengyu Ma

December 28, 2016

Abstract

We give a novel formal theoretical framework for unsupervised learning with two distinctive characteristics. First, it does not assume any generative model and based on a worst-case performance metric. Second, it is comparative, namely performance is measured with respect to a given hypothesis class. This allows to avoid known computational hardness results and improper algorithms based on convex relaxations. We show how several families of unsupervised learning models, which were previously only analyzed under probabilistic assumptions and are otherwise provably intractable, can be efficiently learned in our framework by convex optimization.

1 Introduction

Unsupervised learning is the task of learning structure from unlabelled examples. Informally, the main goal of unsupervised learning is to extract structure from the data in a way that will enable efficient learning from future labelled examples for potentially numerous independent tasks.

It is useful to recall the Probably Approximately Correct (PAC) learning theory for supervised learning [Val84], based on Vapnik’s statistical learning theory [Vap98]. In PAC learning, the learning can access labelled examples from an unknown distribution. On the basis of these examples, the learner constructs a hypothesis that generalizes to unseen data. A concept is said to be learnable with respect to a hypothesis class if there exists an (efficient) algorithm that outputs a generalizing hypothesis with high probability after observing polynomially many examples in terms of the input representation.

The great achievements of PAC learning that made it successful are its generality and algorithmic applicability: PAC learning does not restrict the input domain in any way, and thus allows very general learning, without generative or distributional assumptions on the world. Another important feature is the restriction to specific hypothesis classes, without which there are simple impossibility results such as the “no free lunch” theorem. This allows *comparative* and *improper* learning of computationally-hard concepts.

The latter is a very important point which is often understated. Consider the example of sparse regression, which is a canonical problem in high dimensional statistics. Fitting the best sparse vector to linear prediction is an NP-hard problem [Nat95]. However, this does not prohibit improper learning, since we can use a ℓ_1 convex relaxation for the sparse vectors (famously known as LASSO [Tib96]).

Unsupervised learning, on the other hand, while extremely applicative and well-studied, has not seen such an inclusive theory. The most common approaches, such as restricted Boltzmann machines, topic models, dictionary learning, principal component analysis and metric clustering, are based almost entirely on generative assumptions about the world. This is a strong restriction which makes it very hard to analyze such approaches in scenarios for which the assumptions do not hold. A more discriminative approach is based on compression, such as the Minimum Description Length criterion. This approach gives rise to provably intractable problems and doesn’t allow improper learning.

Main results. We start by proposing a rigorous framework for unsupervised learning which allows data-dependent, comparative learning without generative assumptions. It is general enough to encompass previous methods such as PCA, dictionary learning and topic models. Our main contribution are optimization-based relaxations and efficient algorithms that are shown to improperly probably learn previous models, specifically:

1. We consider the classes of hypothesis known as dictionary learning. We give a more general hypothesis class which encompasses and generalizes it according to our definitions. We proceed to give novel polynomial-time algorithms for learning the broader class. These algorithms are based on new techniques in unsupervised learning, namely sum-of-squares convex relaxations.

As far as we know, this is the first result for efficient improper learning of dictionaries without generative assumptions. Moreover, our result handles polynomially over-complete dictionaries, while previous works [AGMM15, BKS15a] apply to at most constant factor over-completeness.

2. We give efficient algorithms for learning a new hypothesis class which we call spectral autoencoders. We show that this class generalizes, according to our definitions, the class of PCA (principal component analysis) and its kernel extensions.

Structure of this paper. In the following chapter we a non-generative, distribution-dependent definition for unsupervised learning which mirrors that of PAC learning for supervised learning. We then proceed to an illustrative example and show how Principal Component Analysis can be formally learned in this setting. The same section also gives a much more general class of hypothesis for unsupervised learning which we call polynomial spectral decoding, and show how they can be efficiently learned in our framework using convex optimization. Finally, we get to our main contribution: a convex optimization based methodology for improper learning a wide class of hypothesis, including dictionary learning.

1.1 Previous work

The vast majority of work on unsupervised learning, both theoretical as well as applicative, focuses on generative models. These include topic models [BNJ03], dictionary learning [DH06], Deep Boltzmann Machines and deep belief networks [Sal09] and many more. Many times these models entail non-convex optimization problems that are provably NP-hard to solve in the worst-case.

A recent line of work in theoretical machine learning attempts to give efficient algorithms for these models with provable guarantees. Such algorithms were given for topic models [AGM12], dictionary learning [AGM13, AGMM15], mixtures of gaussians and hidden Markov models [HK13, AGH⁺14] and more. However, these works retain, and at times even enhance, the probabilistic generative assumptions of the underlying model. Perhaps the most widely used unsupervised learning methods are clustering algorithms such as k-means, k-medians and principal component analysis (PCA), though these lack generalization guarantees. An axiomatic approach to clustering was initiated by Kleinberg [Kle03] and pursued further in [BDA09]. A discriminative generalization-based approach for clustering was undertaken in [BBV08] within the model of similarity-based clustering.

Another approach from the information theory literature studies with online lossless compression. The relationship between compression and machine learning goes back to the Minimum Description Length criterion [Ris78]. More recent work in information theory gives online algorithms that attain optimal compression, mostly for finite alphabets [ADJ⁺13, OSZ04]. For infinite alphabets, which are the main object of study for unsupervised learning of signals such as images, there are known impossibility results [JOS05]. This connection to compression was recently further advanced, mostly in the context of textual data [PVMH13].

In terms of lossy compression, Rate Distortion Theory (RDT) [Ber71, CT06] is intimately related to our definitions, as a framework for finding lossy compression with minimal distortion (which would

correspond to reconstruction error in our terminology). Our learnability definition can be seen of an extension of RDT to allow improper learning and generalization error bounds. Another learning framework derived from lossy compression is the information bottleneck criterion [TPB00], and its learning theoretic extensions [SST08]. The latter framework assumes an additional feedback signal, and thus is not purely unsupervised.

The downside of the information-theoretic approaches is that worst-case competitive compression is provably computationally hard under cryptographic assumptions. In contrast, our compression-based approach is based on learning a restriction to a specific hypothesis class, much like PAC-learning. This circumvents the impossibility results and allows for improper learning.

2 A formal framework for unsupervised learning

The basis constructs in an unsupervised learning setting are:

1. Instance domain \mathcal{X} , such as images, text documents, etc. Target space, or range, \mathcal{Y} . We usually think of $\mathcal{X} = \mathbb{R}^d, \mathcal{Y} = \mathbb{R}^k$ with $d \gg k$. (Alternatively, \mathcal{Y} can be all sparse vectors in a larger space.)
2. An unknown, arbitrary distribution \mathcal{D} on domain \mathcal{X} .
3. A hypothesis class of decoding and encoding pairs,

$$\mathcal{H} \subseteq \{(h, g) \in \{\mathcal{X} \mapsto \mathcal{Y}\} \times \{\mathcal{Y} \mapsto \mathcal{X}\}\},$$

where h is the encoding hypothesis and g is the decoding hypothesis.

4. A loss function $\ell : \mathcal{H} \times \mathcal{X} \mapsto \mathbb{R}_{\geq 0}$ that measures the reconstruction error,

$$\ell((g, h), x).$$

For example, a natural choice is the ℓ_2 -loss $\ell((g, h), x) = \|g(h(x)) - x\|_2^2$. The rationale here is to learn structure without significantly compromising supervised learning for *arbitrary* future tasks. Near-perfect reconstruction is sufficient as formally proved in Appendix A.1. Without generative assumptions, it can be seen that near-perfect reconstruction is also necessary.

For convenience of notation, we use f as a shorthand for $(h, g) \in \mathcal{H}$, a member of the hypothesis class \mathcal{H} . Denote the generalization ability of an unsupervised learning algorithm with respect to a distribution \mathcal{D} as

$$\text{loss}_{\mathcal{D}}(f) = \mathbb{E}_{x \sim \mathcal{D}} [\ell(f, x)].$$

We can now define the main object of study: unsupervised learning with respect to a given hypothesis class. The definition is parameterized by real numbers: the first is the encoding length (measured in bits) of the hypothesis class. The second is the bias, or additional error compared to the best hypothesis. Both parameters are necessary to allow improper learning.

Definition 2.1. We say that instance \mathcal{D}, \mathcal{X} is (k, γ) - \mathcal{C} -learnable with respect to hypothesis class \mathcal{H} if exists an algorithm that for every $\delta, \varepsilon > 0$, after seeing $m(\varepsilon, \delta) = \text{poly}(1/\varepsilon, \log(1/\delta), d)$ examples, returns an encoding and decoding pair (h, g) (not necessarily from \mathcal{H}) such that:

1. with probability at least $1 - \delta$, $\text{loss}_{\mathcal{D}}((h, g)) \leq \min_{(h, g) \in \mathcal{H}} \text{loss}_{\mathcal{D}}((h, g)) + \varepsilon + \gamma$.
2. $h(x)$ has an explicit representation with length at most k bits.

For convenience we typically encode into real numbers instead of bits. Real encoding can often (though not in the worst case) be trivially transformed to be binary with a loss of logarithmic factor.

Following PAC learning theory, we can use uniform convergence to bound the generalization error of the empirical risk minimizer (ERM). Define the empirical loss for a given sample $S \sim \mathcal{D}^m$ as

$$\text{loss}_S(f) = \frac{1}{m} \cdot \sum_{x \in S} \ell(f, x)$$

Define the ERM hypothesis for a given sample $S \sim \mathcal{D}^m$ as

$$\hat{f}_{ERM} = \arg \min_{f \in \mathcal{H}} \text{loss}_S(f).$$

For a hypothesis class \mathcal{H} , a loss function ℓ and a set of m samples $S \sim \mathcal{D}^m$, define the empirical Rademacher complexity of \mathcal{H} with respect to ℓ and S as,¹

$$\mathcal{R}_{S,\ell}(\mathcal{H}) = \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{f \in \mathcal{H}} \frac{1}{m} \sum_{x \in S} \sigma_i \ell(f, x) \right]$$

Let the Rademacher complexity of \mathcal{H} with respect to distribution \mathcal{D} and loss ℓ as $\mathcal{R}_m(\mathcal{H}) = \mathbb{E}_{S \sim \mathcal{D}^m} [\mathcal{R}_{S,\ell}(\mathcal{H})]$. When it's clear from the context, we will omit the subscript ℓ .

We can now state and apply standard generalization error results. The proof of following theorem is almost identical to [MRT12, Theorem 3.1]. For completeness we provide a proof in Appendix A.

Theorem 2.1. For any $\delta > 0$, with probability $1 - \delta$, the generalization error of the ERM hypothesis is bounded by:

$$\text{loss}_{\mathcal{D}}(\hat{f}_{ERM}) \leq \min_{f \in \mathcal{H}} \text{loss}_{\mathcal{D}}(f) + 6\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{4 \log \frac{1}{\delta}}{2m}}$$

An immediate corollary of the theorem is that as long as the Rademacher complexity of a hypothesis class approaches zero as the number of examples goes to infinity, it can be C learned by an *inefficient* algorithm that optimizes over the hypothesis class by enumeration and outputs an best hypothesis with encoding length k and bias $\gamma = 0$. Not surprisingly such optimization is often intractable and hences the main challenge is to design efficient algorithms. As we will see in later sections, we often need to trade the encoding length and bias slightly for computational efficiency.

Notation. For every vector $z \in \mathbb{R}^{d_1} \otimes \mathbb{R}^{d_2}$, we can view it as a matrix of dimension $d_1 \times d_2$, which is denoted as $\mathcal{M}(z)$. Therefore in this notation, $\mathcal{M}(u \otimes v) = uv^\top$.

Let $v_{\max}(\cdot) : (\mathbb{R}^d)^{\otimes 2} \rightarrow \mathbb{R}^d$ be the function that compute the top right-singular vector of some vector in $(\mathbb{R}^d)^{\otimes 2}$ viewed as a matrix. That is, for $z \in (\mathbb{R}^d)^{\otimes 2}$, then $v_{\max}(z)$ denotes the top right-singular vector of $\mathcal{M}(z)$. We also overload the notation v_{\max} for generalized eigenvectors of higher order tensors. For $T \in (\mathbb{R}^d)^{\otimes \ell}$, let $v_{\max}(T) = \arg \max_{\|x\| \leq 1} T(x, x, \dots, x)$ where $T(\cdot)$ denotes the multi-linear form defined by tensor T .

We use $\|A\|_{\ell_p \rightarrow \ell_q}$ to denote the induced operator norm of A from ℓ_p space to ℓ_q space. For simplicity, we also define $|A|_1 = \|A\|_{\ell_\infty \rightarrow \ell_1} = \sum_{ij} |A_{ij}|$, $|A|_\infty = \|A\|_{\ell_1 \rightarrow \ell_\infty} = \max_{ij} |A_{ij}|$. We note that $\|A\|_{\ell_1 \rightarrow \ell_1}$ is the max column ℓ_1 norm of A , and $\|A\|_{\ell_1 \rightarrow \ell_2}$ is the largest column ℓ_2 norm of A .

¹Technically, this is the Rademacher complexity of the class of functions $\ell \circ \mathcal{H}$. However, since ℓ is usually fixed for certain problem, we emphasize in the definition more the dependency on \mathcal{H} .

3 Spectral autoencoders: unsupervised learning of algebraic manifolds

3.1 Algebraic manifolds

The goal of the spectral autoencoder hypothesis class we define henceforth is to learn the representation of data that lies on a low-dimensional algebraic variety/manifolds. The linear variety, or linear manifold, defined by the roots of linear equations, is simply a linear subspace. If the data resides in a linear subspace, or close enough to it, then PCA is effective at learning its succinct representation.

One extension of the linear manifolds is the set of roots of low-degree polynomial equations. Formally, let k, s be integers and let $c_1, \dots, c_{d^s - k} \in \mathbb{R}^{d^s}$ be a set of vectors in d^s dimension, and consider the algebraic variety

$$\mathcal{M} = \left\{ x \in \mathbb{R}^d : \forall i \in [d^s - k], \langle c_i, x^{\otimes s} \rangle = 0 \right\}.$$

Observe that here each constraint $\langle c_i, x^{\otimes s} \rangle$ is a degree- s polynomial over variables x , and when $s = 1$ the variety \mathcal{M} becomes a linear subspace. Let $a_1, \dots, a_k \in \mathbb{R}^{d^s}$ be a basis of the subspaces orthogonal to all of $c_1, \dots, c_{d^s - k}$, and let $A \in \mathbb{R}^{k \times d^s}$ contains a_i as rows. Then we have that given $x \in \mathcal{M}$, the encoding

$$y = Ax^{\otimes s}$$

pins down all the unknown information regarding x . In fact, for any $x \in \mathcal{M}$, we have $A^\top Ax^{\otimes s} = x^{\otimes s}$ and therefore x is decodable from y . The argument can also be extended to the situation when the data point is close to \mathcal{M} (according to a metric, as we discuss later). The goal of the rest of the subsections is to learn the encoding matrix A given data points residing close to \mathcal{M} .

3.2 Warm up: PCA and kernel PCA

In this section we illustrate our framework for agnostic unsupervised learning by showing how PCA and kernel PCA can be efficiently learned within our model. The results of this sub-section are not new, and given only for illustrative purposes.

The class of hypothesis corresponding to PCA operates on domain $\mathcal{X} = \mathbb{R}^d$ and range $\mathcal{Y} = \mathbb{R}^k$ for some $k < d$ via linear operators. In kernel PCA, the encoding linear operator applies to the s -th tensor power $x^{\otimes s}$ of the data. That is, the encoding and decoding are parameterized by a linear operator $A \in \mathbb{R}^{k \times d^s}$,

$$\mathcal{H}_{k,s}^{\text{pca}} = \left\{ (h_A, g_A) : h_A(x) = Ax^{\otimes s}, g_A(y) = A^\dagger y \right\},$$

where A^\dagger denotes the pseudo-inverse of A . The natural loss function here is the Euclidean norm,

$$\ell((g, h), x) = \|x^{\otimes s} - g(h(x))\|^2 = \|(I - A^\dagger A)x^{\otimes s}\|^2.$$

Theorem 3.1. For a fixed constant $s \geq 1$, the class $\mathcal{H}_{k,s}^{\text{pca}}$ is efficiently C -learnable with encoding length k and bias $\gamma = 0$.

The proof of the Theorem follows from two simple components: a) finding the ERM among $\mathcal{H}_{k,s}^{\text{pca}}$ can be efficiently solved by taking SVD of covariance matrix of the (lifted) data points. b) The Rademacher complexity of the hypothesis class is bounded by $O(d^s/m)$ for m examples. Thus by Theorem 2.1 the minimizer of ERM generalizes. The full proof is deferred to Appendix B.

3.3 Spectral Autoencoders

In this section we give a much broader set of hypothesis, encompassing PCA and kernel-PCA, and show how to learn them efficiently. Throughout this section we assume that the data is normalized to Euclidean norm 1, and consider the following class of hypothesis which naturally generalizes PCA:

Definition 3.1 (Spectral autoencoder). We define the class $\mathcal{H}_{k,s}^{\text{sa}}$ as the following set of all hypothesis (g, h) ,

$$\mathcal{H}_k^{\text{sa}} = \left\{ (h, g) : \begin{array}{l} h(x) = Ax^{\otimes s}, A \in \mathbb{R}^{k \times d^s} \\ g(y) = v_{\max}(By), B \in \mathbb{R}^{d^s \times k} \end{array} \right\}. \quad (3.1)$$

We note that this notion is more general than kernel PCA: suppose some $(g, h) \in \mathcal{H}_{k,s}^{\text{pca}}$ has reconstruction error ε , namely, $A^\dagger Ax^{\otimes s}$ is ε -close to $x^{\otimes s}$ in Euclidean norm. Then by eigenvector perturbation theorem, we have that $v_{\max}(A^\dagger Ax^{\otimes s})$ also reconstructs x with $O(\varepsilon)$ error, and therefore there exists a PSCA hypothesis with $O(\varepsilon)$ error as well. Vice versa, it's quite possible that for every A , the reconstruction $A^\dagger Ax^{\otimes s}$ is far away from $x^{\otimes s}$ so that kernel PCA doesn't apply, but with spectral decoding we can still reconstruct x from $v_{\max}(A^\dagger Ax^{\otimes s})$ since the top eigenvector of $A^\dagger Ax^{\otimes s}$ is close x .

Here the key matter that distinguishes us from kernel PCA is in what metric x needs to be close to the manifold so that it can be reconstructed. Using PCA, the requirement is that x is in Euclidean distance close to \mathcal{M} (which is a subspace), and using kernel PCA $x^{\otimes 2}$ needs to be in Euclidean distance close to the null space of c_i 's. However, Euclidean distances in the original space and lifted space typically are meaningless for high-dimensional data since any two data points are far away with each other in Euclidean distance. The advantage of using spectral autoencoders is that in the lifted space the geometry is measured by spectral norm distance that is much smaller than Euclidean distance (with a potential gap of $d^{1/2}$). The key here is that though the dimension of lifted space is d^2 , the objects of our interests is the set of rank-1 tensors of the form $x^{\otimes 2}$. Therefore, spectral norm distance is a much more effective measure of closeness since it exploits the underlying structure of the lifted data points.

We note that spectral autoencoders relate to vanishing component analysis [LLS⁺13]. When the data is close to an algebraic manifold, spectral autoencoders aim to find the (small number of) essential non-vanishing components in a noise robust manner.

3.4 Learnability of polynomial spectral decoding

For simplicity we focus on the case when $s = 2$. Ideally we would like to learn the best encoding-decoding scheme for any data distribution \mathcal{D} . Though there are technical difficulties to achieve such a general result. A natural attempt would be to optimize the loss function $f(A, B) = \|g(h(x)) - x\|^2 = \|x - v_{\max}(BAx^{\otimes 2})\|^2$. Not surprisingly, function f is not a convex function with respect to A, B , and in fact it could be even non-continuous (if not ill-defined)!

Here we make a further realizability assumption that the data distribution \mathcal{D} admits a reasonable encoding and decoding pair with reasonable reconstruction error.

Definition 3.2. We say a data distribution \mathcal{D} is (k, ε) -regularly spectral decodable if there exist $A \in \mathbb{R}^{k \times d^2}$ and $B \in \mathbb{R}^{d^2 \times k}$ with $\|BA\|_{\text{op}} \leq \tau$ such that for $x \sim \mathcal{D}$, with probability 1, the encoding $y = Ax^{\otimes 2}$ satisfies that

$$\mathcal{M}(By) = \mathcal{M}(BAx^{\otimes 2}) = xx^\top + E, \quad (3.2)$$

where $\|E\|_{\text{op}} \leq \varepsilon$. Here $\tau \geq 1$ is treated as a fixed constant globally.

To interpret the definition, we observe that if data distribution \mathcal{D} is (k, ε) -regularly spectrally decodable, then by equation (3.2) and Wedin's theorem (see e.g. [Vu11]) on the robustness of eigenvector to perturbation, $\mathcal{M}(By)$ has top eigenvector² that is $O(\varepsilon)$ -close to x itself. Therefore, definition 3.2 is a

²Or right singular vector when $\mathcal{M}(By)$ is not symmetric

sufficient condition for the spectral decoding algorithm $v_{\max}(By)$ to return x approximately, though it might be not necessary. Moreover, this condition partially addresses the non-continuity issue of using objective $f(A, B) = \|x - v_{\max}(BAx^{\otimes 2})\|^2$, while $f(A, B)$ remains (highly) non-convex. We resolve this issue by using a convex surrogate.

Our main result concerning the learnability of the aforementioned hypothesis class is:

Theorem 3.2. The hypothesis class $\mathcal{H}_{k,2}^{\text{sa}}$ is C -learnable with encoding length $O(\tau^4 k^4 / \delta^4)$ and bias δ with respect to (k, ε) -regular distributions in polynomial time.

Our approach towards finding an encoding and decoding matrices A, B is to optimize the objective,

$$\begin{aligned} \text{minimize } f(R) &= \mathbb{E} \left[\left\| Rx^{\otimes 2} - x^{\otimes 2} \right\|_{\text{op}} \right] \\ \text{s.t. } \|R\|_{S_1} &\leq \tau k \end{aligned} \quad (3.3)$$

where $\|\cdot\|_{S_1}$ denotes the Schatten 1-norm³.

Suppose \mathcal{D} is (k, ε) -regularly decodable, and suppose h_A and g_B are the corresponding encoding and decoding function. Then we see that $R = AB$ will satisfy that R has rank at most k and $f(R) \leq \varepsilon$. On the other hand, suppose one obtains some R of rank k' such that $f(R) \leq \delta$, then we can produce h_A and g_B with $O(\delta)$ reconstruction simply by choosing $A \in \mathbb{R}^{k' \times d^2}$ and $B \in \mathbb{R}^{d^2 \times k'}$ such that $R = AB$.

We use (non-smooth) Frank-Wolfe to solve objective (3.3), which in particular returns a low-rank solution. We defer the proof of Theorem 3.2 to the Appendix B.1.

With a slightly stronger assumptions on the data distribution \mathcal{D} , we can reduce the length of the code to $O(k^2/\varepsilon^2)$ from $O(k^4/\varepsilon^4)$. See details in Appendix C.

4 A family of optimization encodings and efficient dictionary learning

In this section we give efficient algorithms for learning a family of unsupervised learning algorithms commonly known as "dictionary learning". In contrast to previous approaches, we do not construct an actual "dictionary", but rather improperly learn a comparable encoding via convex relaxations.

We consider a different family of codes which is motivated by matrix-based unsupervised learning models such as topic-models, dictionary learning and PCA. This family is described by a matrix $A \in \mathbb{R}^{d \times r}$ which has low complexity according to a certain norm $\|\cdot\|_\alpha$, that is, $\|A\|_\alpha \leq c_\alpha$. We can parametrize a family of hypothesis \mathcal{H} according to these matrices, and define an encoding-decoding pair according to

$$h_A(x) = \arg \min_{\|y\|_\beta \leq k} \frac{1}{d} |x - Ay|_1, \quad g_A(y) = Ay$$

We choose ℓ_1 norm to measure the error mostly for convenience, though it can be quite flexible. The different norms $\|\cdot\|_\alpha, \|\cdot\|_\beta$ over A and y give rise to different learning models that have been considered before. For example, if these are Euclidean norms, then we get PCA. If $\|\cdot\|_\alpha$ is the max column ℓ_2 or ℓ_∞ norm and $\|\cdot\|_\beta$ is the ℓ_0 norm, then this corresponds to dictionary learning (more details in the next section).

The optimal hypothesis in terms of reconstruction error is given by

$$A^* = \arg \min_{\|A\|_\alpha \leq c_\alpha} \mathbb{E}_{x \sim \mathcal{D}} \left[\frac{1}{d} |x - g_A(h_A(x))|_1 \right] = \arg \min_{\|A\|_\alpha \leq c_\alpha} \mathbb{E}_{x \sim \mathcal{D}} \left[\min_{y \in \mathbb{R}^r: \|y\|_\beta \leq k} \frac{1}{d} |x - Ay|_1 \right].$$

The loss function can be generalized to other norms, e.g., squared ℓ_2 loss, without any essential change in the analysis. Notice that this optimization objective derived from reconstruction error is

³Also known as nuclear norm or trace norm

identically the one used in the literature of dictionary learning. This can be seen as another justification for the definition of unsupervised learning as minimizing reconstruction error subject to compression constraints.

The optimization problem above is notoriously hard computationally, and significant algorithmic and heuristic literature attempted to give efficient algorithms under various distributional assumptions (see [AGM13, AGMM15, AEB05] and the references therein). Our approach below circumvents this computational hardness by convex relaxations that result in learning a different creature, albeit with comparable compression and reconstruction objective.

4.1 Improper dictionary learning: overview

We assume the max column ℓ_∞ norm of A is at most 1 and the ℓ_1 norm of y is assumed to be at most k . This is a more general setting than the random dictionaries (up to a re-scaling) that previous works [AGM13, AGMM15] studied.⁴ In this case, the magnitude of each entry of x is on the order of \sqrt{k} if y has k random ± 1 entries. We think of our target error per entry as much smaller than 1⁵. We consider $\mathcal{H}_{\text{dict}}^k$ that are parametrized by the dictionary matrix $A = \mathbb{R}^{d \times r}$,

$$\mathcal{H}_k^{\text{dict}} = \left\{ (h_A, g_A) : A \in \mathbb{R}^{d \times r}, \|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1 \right\},$$

where $h_A(x) = \arg \min_{\|y\|_1 \leq k} |x - Ay|_1$, $g_A(y) = Ay$

Here we allow r to be larger than d , the case that is often called over-complete dictionary. The choice of the loss can be replaced by ℓ_2 loss (or other Lipschitz loss) without any additional efforts, though for simplicity we stick to ℓ_1 loss. Define A^* to be the the best dictionary under the model and ε^* to be the optimal error,

$$A^* = \arg \min_{\|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1} \mathbb{E}_{x \sim \mathcal{D}} \left[\min_{y \in \mathbb{R}^r: \|y\|_1 \leq k} |x - Ay|_1 \right] \quad (4.1)$$

$$\varepsilon^* = \mathbb{E}_{x \sim \mathcal{D}} \left[\frac{1}{d} \cdot |x - g_{A^*}(h_{A^*}(x))|_1 \right].$$

Algorithm 1 group encoding/decoding for improper dictionary learning

Inputs: N data points $X \in \mathbb{R}^{d \times N} \sim \mathcal{D}^N$. Convex set Q . Sampling probability ρ .

1. **Group encoding:** Compute

$$Z = \arg \min_{C \in Q} |X - C|_1, \quad (4.2)$$

and let

$$Y = h(X) = P_\Omega(Z),$$

where $P_\Omega(B)$ is a random sampling of B where each entry is picked with probability ρ .

2. **Group decoding:** Compute

$$g(Y) = \arg \min_{C \in Q} |P_\Omega(C) - Y|_1. \quad (4.3)$$

Theorem 4.1. For any $\delta > 0, p \geq 1$, the hypothesis class $\mathcal{H}_k^{\text{dict}}$ is C -learnable (by Algorithm 2) with encoding length $\tilde{O}(k^2 r^{1/p} / \delta^2)$, bias $\delta + O(\varepsilon^*)$ and sample complexity $d^{O(p)}$ in time $n^{O(p^2)}$

⁴The assumption can be relaxed to that A has ℓ_∞ norm at most k and ℓ_2 -norm at most \sqrt{d} straightforwardly.

⁵We are conservative in the scaling of the error here. Error much smaller than \sqrt{k} is already meaningful.

We note that here r can be potentially much larger than d since by choosing a large constant p the overhead caused by r can be negligible. Since the average size of the entries is \sqrt{k} , therefore we can get the bias δ smaller than average size of the entries with code length roughly $\approx k$.

The proof of Theorem 4.1 is deferred to Section 5.6. To demonstrate the key intuition and technique behind it, in the rest of the section we consider a simpler algorithm that achieves a *weaker* goal: Algorithm 1 encodes *multiple* examples into some codes with the matching average encoding length $\tilde{O}(k^2 r^{1/p} / \delta^2)$, and these examples can be decoded from the codes together with reconstruction error $\varepsilon^* + \delta$. Next, we outline the analysis of Algorithm 1, and we will show later that one can reduce the problem of encoding a single examples to the problem of encoding multiple examples together.

Here we overload the notation $g_{A^*}(h_{A^*}(\cdot))$ so that $g_{A^*}(h_{A^*}(X))$ denotes the collection of all the $g_{A^*}(h_{A^*}(x_j))$ where x_j is the j -th column of X . Algorithm 1 assumes that there exists a convex set $Q \subset \mathbb{R}^{d \times N}$ such that

$$\left\{ g_{A^*}(h_{A^*}(X)) : X \in \mathbb{R}^{d \times N} \right\} \subset \{AY : \|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1, \|Y\|_{\ell_1 \rightarrow \ell_1} \leq k\} \subset Q. \quad (4.4)$$

That is, Q is a convex relaxation of the group of reconstructions allowed in the class $\mathcal{H}^{\text{dict}}$. Algorithm 1 first uses convex programming to denoise the data X into a clean version Z , which belongs to the set Q . If the set Q has low complexity, then simple random sampling of $Z \in Q$ serves as a good encoding.

The following Lemma shows that if Q has low complexity in terms of sampling Rademacher width, then Algorithm 1 will give a good group encoding and decoding scheme.

Lemma 4.2. Suppose convex $Q \subset \mathbb{R}^{d \times N}$ satisfies condition (4.4). Then, Algorithm 1 gives a group encoding and decoding pair such that with probability $1 - \delta$, the average reconstruction error is bounded by $\varepsilon^* + O(\sqrt{\text{SRW}_m(Q)}) + O(\sqrt{\log(1/\delta)/m})$ where $m = \rho Nd$ and $\text{SRW}_m(\cdot)$ is the sampling Rademacher width (defined in subsection 5.2), and the average encoding length is $\tilde{O}(\rho d)$.

The proofs here are technically standard: Lemma 4.2 simply follows from Lemma 5.1 and Lemma 5.2 in Section 5. Lemma 5.1 shows that the difference between Z and X is comparable to ε^* , which is a direct consequence of the optimization over a large set Q that contains optimal reconstruction. Lemma 5.2 shows that the sampling procedure doesn't lose too much information given a denoised version of the data is already observed, and therefore, one can reconstruct Z from Y .

The novelty here is to use these two steps together to denoise and achieve a short encoding. The typical bottleneck of applying convex relaxation on matrix factorization based problem (or any other problem) is the difficulty of rounding. Here instead of pursuing a rounding algorithm that output the factor A and Y , we look for a convex relaxation that preserves the intrinsic complexity of the set which enables the trivial sampling encoding. It turns out that controlling the width/complexity of the convex relaxation boils down to proving concentration inequalities with sum-of-squares (SoS) proofs, which is conceptually easier than rounding.

Therefore, the remaining challenge is to design convex set Q that simultaneously has the following properties

- (a) is a convex relaxation in the sense of satisfying condition (4.4)
- (b) admits an efficient optimization algorithm
- (c) has low complexity (that is, sampling rademacher width $\tilde{O}(N \text{poly}(k))$)

Most of the proofs need to be deferred to Section 5. We give a brief overview: In subsection 5.3 we will design a convex set Q which satisfies condition (4.4) but not efficiently solvable, and in subsection 5.4 we verify that the sampling Rademacher width is $O(Nk \log d)$. In subsection 5.5, we prove that a sum-of-squares relaxation would give a set Q_p^{SoS} which satisfies (a), (b) and (c) approximately. Concretely, we have the following theorem.

Theorem 4.3. For every $p \geq 4$, let $N = d^{c_0 p}$ with a sufficiently large absolute constant c_0 . Then, there exists a convex set $Q_p^{\text{sos}} \subset \mathbb{R}^{d \times N}$ (which is defined in subsection 5.5.2) such that (a) it satisfies condition 4.4; (b) The optimization (4.2) and (4.3) are solvable by semidefinite programming with run-time $n^{O(p^2)}$; (c) the sampling Rademacher width of Q_p^{sos} is bounded by $\sqrt{\mathcal{SRW}_m(Q)} \leq \tilde{O}(k^2 r^{2/p} N/m)$.

We note that these three properties (with Lemma 4.2) imply that Algorithm 1 with $Q = Q_p^{\text{sos}}$ and $\rho = O(k^2 r^{2/p} d^{-1} / \delta^2 \cdot \log d)$ gives a group encoding-decoding pair with average encoding length $O(k^2 r^{2/p} / \delta^2 \cdot \log d)$ and bias δ .

Proof Overview of Theorem 4.3: At a very high level, the proof exploits the duality between sum-of-squares relaxation and sum-of-squares proof system. Suppose w_1, \dots, w_d are variables, then in SoS relaxation an auxiliary variable W_S is introduced for every subset $S \subset [d]$ of size at most s , and valid linear constraints and psd constraint for W_S 's are enforced. By convex duality, intuitively we have that if a polynomial $q(x) = \sum_{|S| \leq s} \alpha_S x_S$ can be written as a sum of squares of polynomial $q(x) = \sum_j r_j(x)^2$, then the corresponding linear form over X_S , $\sum_{|S| \leq s} \alpha_S X_S$ is also nonnegative. Therefore, to certify certain property of a linear form $\sum_{|S| \leq s} \alpha_S X_S$ over X_S , it is sufficient (and also necessary by duality) that the corresponding polynomial admit a sum-of-squares proof.

Here using the idea above, we first prove the Rademacher width of the convex hull of reconstructions $Q_0 = \text{conv} \{Z = AY : \|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1, \|Y\|_{\ell_1 \rightarrow \ell_1} \leq k\}$ using a SoS proof. Then the same proof automatically applies to for the Rademacher width of the convex relaxation (which is essentially a set of statements about linear combinations of the auxiliary variables). We lose a factor of $r^{2/p}$ because SoS proof is not strong enough for us to establish the optimal Rademacher width of Q_0 .

5 Analysis of Improper Dictionary Learning

In this section we give the full proof of the Theorems and Lemmas in Section 4. We start by stating general results on denoising, Rademacher width and factorizable norms, and proceed to give specialized bounds for our setting in section 5.4.

5.1 Guarantees of denoising

In this subsection, we give the guarantees of the error caused by the denoising step. Recall that ε^* is the optimal reconstruction error achievable by the optimal (proper) dictionary (equation (4.1)).

Lemma 5.1. Let Z be defined in equation (4.2). Then we have that

$$\frac{1}{Nd} \mathbb{E}_{X \sim \mathcal{D}^N} [|Z - X|_1] \leq \varepsilon^* \quad (5.1)$$

Proof. Let $Y^* = A^* h_{A^*}(X)$ where $h_{A^*}(X)$ denote the collection of encoding of X using h_{A^*} . Since $Y^* \in \{AY : \|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1, \|Y\|_{\ell_1 \rightarrow \ell_1} \leq k\} \subset Q$, we have that Y^* is a feasible solution of optimization (4.2). Therefore, we have that $\frac{1}{Nd} \mathbb{E} [|Z - X|_1] \leq \frac{1}{Nd} \mathbb{E} [|X - Y^*|_1] = \varepsilon^*$, where the equality is by the definition of ε^* . \square

5.2 Sampling Rademacher Width of a Set

As long as the intrinsic complexity of the set Q is small then we can compress by random sampling. The idea of viewing reconstruction error the test error of a supervised learning problem started with the work of Srebro and Shraibman [SS05], and has been used for other completion problems, e.g., [BM15]. We use the terminology ‘‘Rademacher width’’ instead of ‘‘Rademacher complexity’’ to emphasize that the notation defined below is a property of a set of vectors (instead of that of a hypothesis class).

For any set $W \subset \mathbb{R}^D$, and an integer m , we define its sampling Rademacher width (SRW) as,

$$\mathcal{SRW}_m(W) = \mathbb{E}_{\sigma, \Omega} \left[\frac{1}{m} \sup_{x \in W} \langle x, \sigma \rangle_{\Omega} \right], \quad (5.2)$$

where Ω is random subset of $[D]$ of size m , $\langle a, b \rangle_{\Omega}$ is defined as $\sum_{i \in \Omega} a_i b_i$ and $\sigma \sim \{\pm 1\}^D$.

Lemma 5.2. ([BM15, Theorem 2.4]) With probability at least $1 - \delta$ over the choice of Ω , for any $\tilde{z} \in \mathbb{R}^D$

$$\frac{1}{D} |\tilde{z} - z|_1 \leq \frac{1}{m} |P_{\Omega}(\tilde{z}) - P_{\Omega}(z)|_1 + 2\mathcal{SRW}_m(W) + M \sqrt{\frac{\log(1/\delta)}{m}}.$$

where $M = \sup_{\tilde{z} \in W, i \in [D]} |z_i - \tilde{z}_i|$.

5.3 Factorable norms

In this subsection, we define in general the factorable norms, from which we obtain a convex set Q which satisfies condition (4.4) (see Lemma 5.3).

For any two norms $\|\cdot\|_{\alpha}, \|\cdot\|_{\beta}$ that are defined on matrices of any dimension, we can define the following quantity

$$\Gamma_{\alpha, \beta}(Z) = \inf_{Z=AB} \|A\|_{\alpha} \|B\|_{\beta} \quad (5.3)$$

For any $p, q, s, t \geq 1$, we use $\Gamma_{p, q, s, t}(\cdot)$ to denote the function $\Gamma_{\ell_p \rightarrow \ell_q, \ell_s \rightarrow \ell_t}(\cdot)$. When $p = t$, $\Gamma_{p, q, s, p}(Z)$ is the factorable norm [TJ89, Chapter 13] of matrix Z . In the case when $p = t = 2$, $q = \infty, s = 1$, we have that $\Gamma_{2, \infty, 1, 2}(\cdot)$ is the γ_2 -norm [LMSS07] or max norm [SS05], which has been used for matrix completion.

The following Lemma is the main purpose of this section which shows a construction of a convex set Q that satisfies condition (4.4).

Lemma 5.3. For any $q, t \geq 1$ we have that $\Gamma_{1, q, 1, t}(\cdot)$ is a norm. As a consequence, letting $Q_{1, \infty, 1, 1} = \{C \in \mathbb{R}^{N \times d} : \Gamma_{1, \infty, 1, 1}(C) \leq \sqrt{dk}\}$, we have that $Q_{1, \infty, 1, 1}$ is a convex set and it satisfies condition (4.4).

Towards proving Lemma 5.3, we prove a stronger result that if $p = s = 1$, then $\Gamma_{p, q, s, t}$ is also a norm. This result is parallel to [TJ89, Chapter 13] where the case of $p = t$ is considered.

Theorem 5.4. Suppose that $\|\cdot\|_{\alpha}$ and $\|\cdot\|_{\beta}$ are norms defined on matrices of any dimension such that

1. $\|[A, B]\|_{\alpha} \leq \max\{\|A\|_{\alpha}, \|B\|_{\alpha}\}$
2. $\|\cdot\|_{\beta}$ is invariant with respect to appending a row with zero entries.

Then, $\Gamma_{\alpha, \beta}(\cdot)$ is a norm.

Proof. Non-negativity: We have that $\Gamma_{\alpha, \beta}(Z) \geq 0$ by definition as a product of two norms. Further, $\Gamma_{\alpha, \beta}(Z) = 0$ if and only if $\|A\|_{\alpha} = 0$ or $\|B\|_{\beta} = 0$, which is true if and only if A or B are zero, which means that Z is zero.

Absolute scalability: For any positive t , we have that if $Z = AB$ and $\Gamma_{\alpha, \beta}(Z) \leq \|A\|_{\alpha} \|B\|_{\beta}$, then $tZ = (tA) \cdot B$ and $\Gamma_{\alpha, \beta}(tZ) \leq t\|A\|_{\alpha} \|B\|_{\beta}$. Therefore by definition of $\Gamma_{\alpha, \beta}(Z)$, we get $\Gamma_{\alpha, \beta}(tZ) \leq t\Gamma_{\alpha, \beta}(Z)$.

If we replace Z by tZ and t by $1/t$ we obtain the other direction, namely $\Gamma_{\alpha, \beta}(Z) \leq 1/t \cdot \Gamma_{\alpha, \beta}(tZ)$. Therefore, $\Gamma_{\alpha, \beta}(tZ) = t\Gamma_{\alpha, \beta}(Z)$ for any $t \geq 0$.

Triangle inequality: We next show that $\Gamma_{\alpha, \beta}(Z)$ satisfies triangle inequality, from which the result follows. Let W and Z be two matrices of the same dimension. Suppose A, C satisfy that $Z = AC$ and

$\Gamma_{\alpha,\beta}(Z) = \|A\|_\alpha \|C\|_\beta$. Similarly, suppose $W = BD$, and $\Gamma_{\alpha,\beta}(W) = \|B\|_\alpha \|D\|_\beta$. Therefore, we have that $W + Z = [tA, B] \begin{bmatrix} t^{-1}C \\ D \end{bmatrix}$, and that for any $t > 0$,

$$\begin{aligned} \Gamma_{\alpha,\beta}(W + Z) &\leq \|[tA, B]\|_\alpha \left\| \begin{bmatrix} t^{-1}C \\ D \end{bmatrix} \right\|_\beta && \text{(by definition of } \Gamma_{\alpha,\beta}\text{)} \\ &\leq \|[tA, B]\|_\alpha \left(\left\| \begin{bmatrix} 0 \\ D \end{bmatrix} \right\|_\beta + t^{-1} \left\| \begin{bmatrix} C \\ 0 \end{bmatrix} \right\|_\beta \right) && \text{(by triangle inequality)} \\ &\leq \max\{t\|A\|_\alpha, \|B\|_\alpha\} \left(t^{-1}\|C\|_\beta + \|D\|_\beta \right) && \text{(by assumptions on } \|\cdot\|_\alpha \text{ and } \|\cdot\|_\beta\text{)} \end{aligned}$$

Pick $t = \frac{\|B\|_\alpha}{\|A\|_\alpha}$, we obtain that,

$$\begin{aligned} \Gamma_{\alpha,\beta}(W + Z) &\leq \|A\|_\alpha \|C\|_\beta + \|B\|_\alpha \|D\|_\beta \\ &= \Gamma_{\alpha,\beta}(Z) + \Gamma_{\alpha,\beta}(W). \end{aligned}$$

□

Note that if $\|\cdot\|_\alpha$ is a $\ell_1 \rightarrow \ell_q$ norm, then it's also the max column-wise ℓ_q norm, and therefore it satisfies the condition a) in Theorem 5.4. Moreover, for similar reason, $\|\cdot\|_\beta = \|\cdot\|_{\ell_1 \rightarrow \ell_t}$ satisfies the condition b) in Theorem 5.4. Hence, Lemma 5.3 is a direct consequence of Theorem 5.4. Lemma 5.3 gives a convex set Q that can be potentially used in Algorithm 1.

5.4 Sampling Rademacher width of level set of $\Gamma_{1,\infty,1,1}$

Here we give a Rademacher width bound for the specific set we're interested in, namely the level sets of $\Gamma_{1,\infty,1,1}$, formally,

$$Q_{1,\infty,1,1} = \{C \in \mathbb{R}^{N \times d} : \Gamma_{1,\infty,1,1}(C) \leq k\}.$$

By the definition $Q_{1,\infty,1,1}$ satisfies condition (4.4). See section 5.2 for definition of Rademacher width.

Lemma 5.5. It holds that

$$SRW_m(Q_{1,\infty,1,1}) \leq \tilde{O} \left(\sqrt{\frac{k^2 N}{m}} \right).$$

Proof of Lemma 5.5. Recall the definition of the sample set Ω of coordinates from C , and their multiplication by i.i.d Rademacher variables in section 5.2. Reusing notation, let $\xi = \sigma \odot \Omega$ and we use Q as a shorthand for $Q_{1,\infty,1,1}$. Here \odot means the entry-wise Hadamard product (namely, each coordinate in Ω is multiplied by an independent Rademacher random variable). We have that

$$\begin{aligned} SRW_m(Q) &= \mathbb{E}_\xi \left[\frac{1}{m} \sup_{C \in Q} \langle C, \xi \rangle \right] \\ &= \mathbb{E}_\xi \left[\frac{1}{m} \sup_{\|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1, \|B\|_{\ell_1 \rightarrow \ell_1} \leq k} \langle AB, \xi \rangle \right] && \text{(by definition of } Q\text{)} \\ &= \mathbb{E}_\xi \left[\frac{1}{m} \sup_{\|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1, \|B\|_{\ell_1 \rightarrow \ell_1} \leq k} \langle B, A^\top \xi \rangle \right] \\ &\leq \mathbb{E}_\xi \left[\frac{1}{m} \sup_{\|A\|_{\ell_1 \rightarrow \ell_\infty} \leq 1} k \sum_{i=1}^N \|A^\top \xi_i\|_\infty \right]. && \text{(By } \langle U, V \rangle \leq \left(\sum_{i=1}^N \|U_i\|_\infty \right) \|V\|_{\ell_1 \rightarrow \ell_1}\text{)} \end{aligned}$$

Let $\rho = \frac{m}{dN}$ be the probability of any entry belongs to Ω . Let ξ_i denote the i -th column of ξ , and A_j denotes the j -th column of A . Therefore, each entry of ξ_i has probability $\rho/2$ to be $+1$ and -1 , and has probability $1 - \rho$ of being 0 . By concentration inequality we have that for $\rho \geq \frac{\log r}{d}$, and any fixed A with $\|A\|_{\ell_1 \rightarrow \ell_\infty} = \max \|A_j\|_\infty \leq 1$,

$$\mathbb{E}_{\xi_i}[\|A^\top \xi_i\|_\infty] \leq O(\sqrt{\rho d} \log r \log d). \quad (5.4)$$

Moreover, we have that

$$\text{Var}_{\xi_i}[\|A^\top \xi_i\|_\infty] \leq O(\sqrt{\rho d} \log r \log d). \quad (5.5)$$

Moreover, $\|A^\top \xi_i\|_\infty$ has an sub-exponential tail. (Technically, its ψ_1 -Orlicz norm is bounded by $O(\sqrt{\rho d} \log r \log d)$). Note that the variance of $\sum_{i=1}^N \|A^\top \xi_i\|_\infty$ will decrease as N increases, and therefore for large enough $N = (dr\rho)^{\Omega(1)}$, we will have that with probability $1 - \exp(-\Omega(dr)^{\Omega(1)})$,

$$\sum_{i=1}^N \|A^\top \xi_i\|_\infty \leq O(N \sqrt{\rho d \log r \log d})$$

Therefore, using the standard ε -net covering argument, we obtain that with high probability,

$$\sup_{\|A\|_{\ell_1 \rightarrow \ell_2} \leq \sqrt{d}} \sum_{i=1}^N \|A^\top \xi_i\|_\infty \leq O(N \sqrt{\rho d \log r \log d}).$$

Hence, altogether we have

$$SRW_m(Q) \leq \mathbb{E}_\xi \left[\frac{1}{m} \sup_{\|A\|_{\ell_1 \rightarrow \ell_2} \leq \sqrt{d}} k \|A^\top \xi\|_{\ell_1 \rightarrow \ell_\infty} \right] \leq \tilde{O} \left(\sqrt{\frac{k^2 N}{m}} \right).$$

□

5.5 Convex Relaxation for $\Gamma_{1,\infty,1,1}$ norm

5.5.1 Sum-of-squares relaxation

Here we will only briefly introduce the basic ideas of Sum-of-Squares (Lasserre) relaxation [Par00, Las01] that will be used for this paper. We refer readers to the extensive study [Las15, Lau09, BS14] for detailed discussions of sum of squares proofs and their applications to algorithm design. Recently, there has been a popular line of research on applications of sum-of-squares algorithms to machine learning problems [BKS15b, BKS14, BM16, MW15, GM15, HSS15, HSS16, MSS16]. Here our technique in the next subsection is most related to that of [BM16], with the main difference that we deal with ℓ_1 norm constraints that are not typically within the SoS framework.

Let $\mathbb{R}[x]_d$ denote the set of all real polynomials of degree at most d with n variables x_1, \dots, x_n . We start by defining the notion of *pseudo-expectation*. The intuition is that the pseudo-expectation behave like the actual expectation of a real probability distribution on squares of polynomials.

Definition 5.1 (pseudo-expectation). A degree- d pseudo-expectation $\tilde{\mathbb{E}}$ is a linear operator that maps $\mathbb{R}[x]_d$ to \mathbb{R} and satisfies $\tilde{\mathbb{E}}(1) = 1$ and $\tilde{\mathbb{E}}(p^2(x)) \geq 0$ for all real polynomials $p(x)$ of degree at most $d/2$.

Definition 5.2. Given a set of polynomial equations $\mathcal{A} = \{q_1(x) = 0, \dots, q_n(x) = 0\}$, we say degree- d pseudo-expectation $\tilde{\mathbb{E}}$ satisfies constraints \mathcal{A} if $\tilde{\mathbb{E}}[q_i(x)r(x)] = 0$ for every i and $r(x)$ such that $\deg(q_i r) \leq d$.

One can optimize over the set of pseudo-expectations that satisfy \mathcal{A} in $n^{O(d)}$ time by the following semidefinite program:

$$\begin{aligned} \text{Variables} \quad & \tilde{\mathbb{E}}[x^S] && \forall S : |S| \leq d \\ \text{Subject to} \quad & \tilde{\mathbb{E}}[q_i(x)x^K] = 0 && \forall i, K : |K| + \deg(q_i) \leq d \\ & \tilde{\mathbb{E}}[x^{\otimes d/2}(x^{\otimes d/2})^\top] \succeq 0 \end{aligned}$$

Definition 5.3 (SoS proof of degree d). For a set of constraints $\mathcal{A} = \{q_1(x) = 0, \dots, q_n(x) = 0\}$, and an integer d , we write

$$\mathcal{A} \vdash_d p(x) \geq q(x)$$

if there exists polynomials $h_i(x)$ for $i = 0, 1, \dots, \ell$ and $g_j(x)$ for $j = 1, \dots, t$ such that $\deg(h_i) \leq d/2$ and $\deg(g_j r_j) \leq d$ that satisfy

$$p(x) - q(x) = \sum_{i=1}^{\ell} h_i(x)^2 + \sum_{j=1}^t r_j(x)g_j(x).$$

We will drop the subscript d when it is clear from the context.

The following fact follows from the definitions above but will be useful throughout the proof.

Proposition 5.6. Suppose $\mathcal{A} \vdash_d p(x) \geq q(x)$. Then for any degree- d pseudo-expectation $\tilde{\mathbb{E}}$ that satisfies \mathcal{A} , we have $\tilde{\mathbb{E}}[p(x)] \geq \tilde{\mathbb{E}}[q(x)]$.

5.5.2 Relaxation for $\Gamma_{1,\infty,1,1}$ norm

In this section, we introduce convex relaxations for the $\Gamma_{1,\infty,1,1}$ norm. For convenience, let p be a power of 2. Let A and B be formal variables of dimension $d \times r$ and $r \times N$ in this section. We introduce more formal variables for the relaxation. Let b be formal variables of dimension $r \times N$. We consider the following set of polynomial constraints over formal variables A, B, b :

$$\mathcal{A} = \left\{ \forall i, j, B_{ij} = b_{ij}^{p-1}, \sum_{\ell=1}^r b_{\ell j}^p \leq k^{p/(p-1)}, \forall i, k, A_{ik}^2 \leq 1 \right\}.$$

For any *real* matrix $C \in \mathbb{R}^{d \times N}$, we define

$$\mathcal{A}(C) = \{C = AB\} \cup \mathcal{A}.$$

We define our convex relaxation for $Q_{1,\infty,1,1}$ as follows,

$$Q_p^{\text{sos}} = \{C \in \mathbb{R}^{d \times N} : \exists \text{degree-}O(p^2) \text{ pseudo-expectation } \tilde{\mathbb{E}} \text{ that satisfies } \mathcal{A}(C)\} \quad (5.6)$$

Lemma 5.7. For any $p \geq 4$, we have

$$Q_{1,\infty,1,1} \subset Q_p^{\text{sos}}$$

and therefore Q_p^{sos} satisfies condition (4.4).

Proof. Suppose $C \in Q_{1,\infty,1,1}$. Then by definition of the $\Gamma_{1,\infty,1,1}$ -norm (equation (5.3)), we have that there exists matrices U, V such that $C = UV$ and $U_{ij}^2 \leq 1$ and $\|V\|_{\ell_1 \rightarrow \ell_1} \leq k$. Now construct $v_{ij} = V_{ij}^{1/(p-1)}$. We have $\sum_{\ell=1}^r b_{\ell j}^p \leq \left(\sum_{\ell=1}^r v_{\ell j}^{p-1}\right)^{p/(p-1)} \leq k^{p/(p-1)}$. Therefore, Then we have that $A = U, B = V, b = v$ satisfies the constraint (5.6). Then the trivial pseudo-expectation operator $\tilde{\mathbb{E}}[p(A, B, b)] = p(U, V, v)$ satisfies $\mathcal{A}(C)$ and therefore $C \in Q_p^{\text{sos}}$ by definition. \square

Theorem 5.8. Suppose $N = d^{c_0 p}$ for large enough absolute constant c_0 . Then the sampling Rademacher complexity of Q_p^{SoS} is bounded by,

$$\mathcal{SRW}_m(Q_p^{\text{SoS}}) \leq O\left(\sqrt{\frac{p^2 N k^2 r^{2/p} \log d}{m}}\right).$$

The proof of Theorem 5.8 finally boils down to prove certain empirical process statement with SoS proofs. We start with the following two lemmas.

Lemma 5.9. Suppose $N = d^{c_0 p}$ for larger enough constant c_0 , and let $\xi = \sigma \odot \Omega$ where σ and Ω are defined in Section 5.2. Then, we have

$$\mathcal{A} \vdash \langle AB, \xi \rangle^p \leq N^{p-1} k^p \sum_{i=1}^N \|A^\top \xi_i\|_p^p$$

Proof. Let ξ_i be the i -th column of ξ . We have that

$$\begin{aligned} \mathcal{A} \vdash \langle AB, \xi \rangle^p &= \langle A^\top \xi, B \rangle^p = \left(\sum_{i=1}^N \langle A^\top \xi, B_i \rangle \right)^p \\ &\leq N^{p-1} \left(\sum_{i=1}^N \langle A^\top \xi_i, B_i \rangle^p \right) && \text{(since } \vdash (\frac{1}{N} \sum_{i \in [N]} \alpha_i)^p \leq \frac{1}{N} \sum_{i \in [N]} \alpha_i^p) \\ &= N^{p-1} \sum_{i=1}^N \langle A^\top \xi_i, b_i^{\odot p-1} \rangle^p && \text{(by } B_{ij} = b_{ij}^{p-1}) \\ &\leq N^{p-1} \sum_{i=1}^N \|A^\top \xi_i\|_p^p \|b_i\|_p^{(p-1)p} && \text{(by } \vdash (\sum_j a_j b_j^{p-1})^p \leq (\sum_i a_i^p) (\sum b_i^p)^{p-1}; \text{ see Lemma D.1)} \\ &\leq N^{p-1} k^p \sum_{i=1}^N \|A^\top \xi_i\|_p^p && \text{(by constraint } \sum_{\ell=1}^r b_{\ell j}^p \leq k^{p/(p-1)}) \end{aligned}$$

□

Lemma 5.10. In the setting of Lemma 5.9, let $\rho = m/(Nd)$. Let $x = (x_1, \dots, x_d)$ be indeterminates and let $\mathcal{B} = \{x_1^2 \leq 1, \dots, x_d^2 \leq 1\}$. Suppose $\rho \geq 1/d$. With probability at least $1 - \exp(-\Omega(d))$ over the choice of ξ , there exists a SoS proof,

$$\mathcal{B} \vdash \|\xi^\top x\|_p^p \leq N \cdot O(\rho d p^2)^{p/2}. \quad (5.7)$$

As an immediate consequence, let ξ_i be the i -th column of ξ , then, with probability at least $1 - \exp(-\Omega(d))$ over the choice of ξ there exists SoS proof,

$$\mathcal{A} \vdash \sum_{i=1}^N \|A^\top \xi_i\|_p^p \leq Nr \cdot O(\rho d p^2)^{p/2}. \quad (5.8)$$

Proof. Recall that p is an even number. We have that

$$\mathcal{B} \vdash \|\xi^\top x\|_p^p = \left(x^{\otimes p/2}\right)^\top \left(\sum_{i=1}^N \xi_i^{\otimes p/2} \left(\xi_i^{\otimes p/2}\right)^\top\right) x^{\otimes p/2}$$

Let $T = \sum_{i=1}^N \xi_i^{\otimes p/2} (\xi_i^{\otimes p/2})^\top \in \mathbb{R}^{d^{p/2} \times d^{p/2}}$. Then, by definition we have that $\mathbb{E} \left[(x^{\otimes p/2})^\top T x^{\otimes p/2} \right] = \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right]$. It follows that

$$\begin{aligned} \mathcal{B} \vdash \|\xi^\top x\|_p^p &= (x^{\otimes p/2})^\top (T - \mathbb{E}[T]) x^{\otimes p/2} + \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right] \\ &\leq \|x\|^p \|T - \mathbb{E}[T]\| + \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right] \quad (\text{by } \vdash y^\top B y \leq \|y\|^2 \|B\|) \\ &\leq d^{p/2} \|T - \mathbb{E}[T]\| + \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right]. \end{aligned}$$

Let ζ have the same distribution as ξ_i . Then we have that

$$\mathbb{E} [\langle \zeta, x \rangle^p] = \sum_{\alpha} t_{\alpha} x^{\alpha}.$$

where $t_{\alpha} = 0$ if x^{α} contains an odd degree. Moreover, let $|\alpha|$ be the number x_i with non-zero exponents that x^{α} has. Then we have $t_{\alpha} \leq \rho^{|\alpha| p}$. Therefore we have that for $\rho \geq 1/d$,

$$\mathcal{B} \vdash \mathbb{E} [\langle \zeta, x \rangle^p] \leq d^{p/2} \rho^{p/2} p^p.$$

It follows that

$$\mathcal{B} \vdash \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right] \leq N \cdot O(\rho d p^2)^{p/2}.$$

Next, we use Bernstein inequality to bound $\|T - \mathbb{E}[T]\|$. We control the variance by

$$\sigma^2 \triangleq \left\| \mathbb{E} \left[\sum_{i=1}^N \|\xi\|^p \xi_i^{\otimes p/2} (\xi_i^{\otimes p/2})^\top \right] \right\| \leq N d^{2p}.$$

Moreover, each summand in the definition of T can be bounded by $\left\| \xi_i^{\otimes p/2} (\xi_i^{\otimes p/2})^\top \right\| \leq d^p$. Note that these bounds are not necessarily tight since they already suffice for our purpose. Therefore, we obtain that with high probability over the chance of ξ , it holds that $\|T - \mathbb{E}[T]\| \leq \sqrt{N d^{O(p)} \log N}$. Therefore for $N \geq d^{\Omega(p)}$, we have that with high probability over the choice of ξ ,

$$\begin{aligned} \mathcal{B} \vdash \|\xi^\top x\|_p^p & \tag{5.9} \\ &\leq d^{p/2} \|T - \mathbb{E}[T]\| + \mathbb{E} \left[\sum_{i=1}^N \langle \xi_i, x \rangle^p \right] \leq N \cdot O(\rho d p^2)^{p/2}. \end{aligned}$$

Finally we show equation (5.7) implies equation (5.8). Let A_i be the i -th column of A . Note that we have $\sum_{i=1}^N \|A^\top \xi_i\|_p^p = \sum_{i=1}^r \|\xi^\top A_i\|_p^p$, since both left-hand side and right-hand side are equal to the p -th power of all the entries of the matrix $A^\top \xi$. Since $\mathcal{A} \vdash \{\forall j, A_{ij}^2 \leq 1\}$, we can invoke the first part of the Lemma and use equation (5.7) to complete the proof. \square

Combining Lemma 5.9 and Lemma 5.10 we obtain the following corollary,

Corollary 5.11. In the setting of Lemma 5.9, with probability $1 - \exp(-\Omega(d))$, we have

$$\mathcal{A} \vdash \langle AB, \xi \rangle^p \leq N^p k^p r \cdot O(\rho d)^{p/2}. \tag{5.10}$$

Now we are ready to prove Theorem 5.8. Essentially the only thing that we need to do is take pseudo-expectation on the both sides of equation (5.10).

Proof of Theorem 5.8. Recall that $\xi = \sigma \odot \Omega$. We have that

$$\begin{aligned}
\mathcal{SRW}_m(Q_p^{\text{sos}})^p &= \mathbb{E}_\xi \left[\frac{1}{m^p} \sup_{C \in Q_d^{\text{sos}}} \langle C, \xi \rangle \right]^p \\
&\leq \mathbb{E}_\xi \left[\frac{1}{m^p} \sup_{C \in Q_d^{\text{sos}}} \langle C, \xi \rangle^p \right] && \text{(by Jensen's inequality)} \\
&= \mathbb{E}_\xi \left[\frac{1}{m^p} \sup_{\tilde{\mathbb{E}} \text{ that satisfies } \mathcal{A}(C)} \tilde{\mathbb{E}}[\langle AB, \xi \rangle] \right] && \text{(by definition of } Q_p^{\text{sos}}) \\
&\leq \mathbb{E} \left[\frac{1}{m^p} \sup_{\tilde{\mathbb{E}} \text{ that satisfies } \mathcal{A}(C)} \tilde{\mathbb{E}}[\langle AB, \xi \rangle] \mid \text{equation (5.8) holds} \right] \\
&\quad + \mathbb{P}[\text{equation (5.8) doesn't hold}] \cdot d^{O(1)} \\
&\leq \frac{1}{m^p} N^p k^p r \cdot O(\rho d p^2)^{p/2} + \exp(-\Omega(d)) d^{O(1)} \\
&\hspace{15em} \text{(by Corollary 5.11 and Proposition 5.6)} \\
&\leq \frac{1}{m^p} N^p k^p r \cdot O(\rho d p^2)^{p/2}.
\end{aligned}$$

Therefore using the fact that $\rho = m/(Nd)$, taking p -th root on the equation above, we obtain,

$$\mathcal{SRW}_m(Q_p^{\text{sos}}) \leq O \left(\sqrt{\frac{p^2 N r^2 / p k^2 \log d}{m}} \right).$$

□

5.6 Proof of Theorem 4.1

In this subsection we prove Theorem 4.1 using the machinery developed in previous subsections. Essentially the idea here is to reduce the problem of encoding one example to the problem of encoding a group of examples. To encode an example x , we draw $N - 1$ fresh examples x_1, \dots, x_{N-1} , and then call the group encoding Algorithm 1.

We also note that the encoding procedure can be simplified by removing the denoising step, and this still allows efficient decoding steps. Since in this case the encoding contains more noise from the data, we prefer the current scheme.

Proof of Theorem 4.1. We will prove that Algorithm 2 with $Q = Q_{\text{sos}}^p$ and $\rho = O(k^2 r^2 / p d^{-1} / \delta^2 \cdot \log d)$ gives the desired encoding/decoding. Indeed with high probability, the encoding length is $\tilde{O}(\rho d) = \tilde{O}(k^2 r^2 / p / \delta^2)$. Next we will show that we obtain good reconstruction in the sense that $\frac{1}{d} \mathbb{E} [|x - g(h(x))|_1] = \frac{1}{d} \mathbb{E} [|x - \tilde{z}|_1] \leq O(\varepsilon^* + \delta)$. Here and throughout the proof, the expectation is over the randomness of x and the randomness of the algorithm (that is, the randomness of x, x_1, \dots, x_{N-1} and the random sampling). First of all, by symmetry, we have that

$$\mathbb{E} [|x - \tilde{z}|_1] = \frac{1}{Nd} \mathbb{E} [|x_1, \dots, x_{N-1}, x] - [\tilde{z}_1, \dots, \tilde{z}_{N-1}, \tilde{z}]|_1]. \quad (5.12)$$

Let $X = [x_1, \dots, x_{N-1}, x]$, and $\tilde{Z} = [\tilde{z}_1, \dots, \tilde{z}_{N-1}, \tilde{z}]$ for simplicity of notations. Let $[Z_{1,i}, z_i] = \arg \min_{[C,c] \in Q} |[X_{1,i}, x_i] - [C, c]|_1$ where $X_{1,i}$ be the samples drawn in the encoding process for x_i .

Algorithm 2 Encoding/decoding for improper dictionary learning

Given: Integer N . Convex set $Q \subset \mathbb{R}^{d \times N}$. Sampling probability ρ .

1. **Encoding: input:** data point $x \in \mathbb{R}^d \sim \mathcal{D}$, **output:** codeword $y = h(x) \in \mathbb{R}^{k \times N}$ Draw $N - 1$ data points $X_1 \in \mathbb{R}^{d \times N-1} \sim \mathcal{D}^{N-1}$.

Denoising step:

$$[Z_1, z] = \arg \min_{[C, c] \in Q} |[X_1, x] - [C, c]|_1, \quad (5.11)$$

Random sampling step:

$$h(x) = P_\Omega(z),$$

where $P_\Omega(z)$ is a random sampling of z where each entry is picked with probability ρ .

2. **Decoding: input:** codeword $y \in \mathbb{R}^{k \times N}$ **output:** reconstruction $g(y) \in \mathbb{R}^{d \times N}$

Take N more samples x_1, \dots, x_N . Encode them to $y_1 = h(x_1), \dots, y_N = h(x_N)$. Then, compute

$$[\tilde{z}_1, \dots, \tilde{z}_{N-1}, \tilde{z}] = \arg \min_{C \in Q} |P_\Omega([C, c]) - [y_1, \dots, y_N, y]|_1.$$

Return $g(y) = \tilde{z}$.

Let $Z = [z_1, \dots, z_{N-1}, z]$. By Lemma 5.1 and the symmetry, we have that for any i , $\frac{1}{d} \mathbb{E} [|z_i - x_i|_1] \leq \frac{1}{dN} \mathbb{E} [|[Z_{1,i}, z_i] - [X_{1,i}, x_i]|_1] \leq \varepsilon^*$. Thus, we have $\frac{1}{dN} \mathbb{E} [|X - Z|_1] \leq \varepsilon^*$.

Let $\hat{Z} \in \mathbb{R}^{d \times N}$ be a helper matrix in the analysis defined as

$$\hat{Z} = \arg \min_{C \in Q} |X - C|_1. \quad (5.13)$$

Then by Lemma 5.1 we have that $\frac{1}{Nd} \mathbb{E} [|\hat{Z} - X|_1] \leq \varepsilon^*$. Then by triangle inequality we obtain that

$$\frac{1}{Nd} \mathbb{E} [|\hat{Z} - Z|_1] \leq 2\varepsilon^*. \quad (5.14)$$

Note that by definition $\hat{Z} \in Q$. Therefore, by Lemma 5.2, we have that

$$\begin{aligned} \frac{1}{Nd} \mathbb{E} [|\hat{Z} - Z|_1] &\leq \frac{1}{m} \mathbb{E} [|P_\Omega(\hat{Z}) - P_\Omega(Z)|_1] + 2SRW_m(W) && \text{(by Lemma 5.2)} \\ &\leq \frac{1}{m} \mathbb{E} [|P_\Omega(\hat{Z}) - P_\Omega(Z)|_1] + 2SRW_m(W) \\ &&& \text{(since } \hat{Z} \text{ is the minimizer, and } \hat{Z} \in Q.) \\ &= \frac{1}{Nd} \mathbb{E} [|\hat{Z} - Z|_1] + 2SRW_m(W) \\ &\leq \varepsilon^* + \delta. && \text{(using equation (5.14) and Theorem 4.3 (c))} \end{aligned}$$

Finally by triangle inequality again we have

$$\frac{1}{Nd} \mathbb{E} [|\tilde{Z} - X|_1] \leq \frac{1}{Nd} \mathbb{E} [|\tilde{Z} - Z|_1] + \frac{1}{Nd} \mathbb{E} [|X - Z|_1] \leq \varepsilon^*. \quad (5.15)$$

Combining equation (5.12) and (5.15) we complete the proof. \square

6 Conclusions

We have defined a new framework for unsupervised learning which replaces generative assumptions by notions of reconstruction error and encoding length. This framework is comparative, and allows learning of particular hypothesis classes with respect to an unknown distribution by other hypothesis classes.

We demonstrate its usefulness by giving new polynomial time algorithms for two unsupervised hypothesis classes. First, we give new polynomial time algorithms for dictionary models in significantly broader range of parameters and assumptions. Another domain is the class of spectral encodings, for which we consider a new class of models that is shown to strictly encompass PCA and kernel-PCA. This new class is capable, in contrast to previous spectral models, learn algebraic manifolds. We give efficient learning algorithms for this class based on convex relaxations.

Acknowledgements

We thank Sanjeev Arora for many illuminating discussions and crucial observations in earlier phases of this work, amongst them that a representation which preserves information for all classifiers requires lossless compression.

References

- [ADJ⁺13] Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Tight bounds for universal compression of large alphabets. In *Proceedings of the 2013 IEEE International Symposium on Information Theory, Istanbul, Turkey, July 7-12, 2013*, pages 2875–2879, 2013.
- [AEB05] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: Design of dictionaries for sparse representation. In *IN: PROCEEDINGS OF SPARS05*, pages 9–12, 2005.
- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *J. Mach. Learn. Res.*, 15(1):2773–2832, January 2014.
- [AGM12] Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models—going beyond svd. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 1–10. IEEE, 2012.
- [AGM13] Sanjeev Arora, Rong Ge, and Ankur Moitra. New algorithms for learning incoherent and overcomplete dictionaries. *arXiv preprint arXiv:1308.6273*, 2013.
- [AGMM15] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra. Simple, efficient, and neural algorithms for sparse coding. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 113–149, 2015.
- [BBV08] Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. A discriminative framework for clustering via similarity functions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, pages 671–680, 2008.
- [BDA09] Shai Ben-David and Margareta Ackerman. Measures of clustering quality: A working set of axioms for clustering. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 121–128. Curran Associates, Inc., 2009.

- [Ber71] Toby Berger. Rate distortion theory: A mathematical basis for data compression. 1971.
- [BKS14] Boaz Barak, Jonathan A. Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *STOC*, pages 31–40, 2014.
- [BKS15a] Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 143–151, 2015.
- [BKS15b] Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC '15, 2015*.
- [BM15] Boaz Barak and Ankur Moitra. Tensor prediction, rademacher complexity and random 3-xor. *CoRR*, abs/1501.06521, 2015.
- [BM16] Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 417–445, 2016.
- [BNJ03] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, March 2003.
- [BS14] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, 2014. To appear.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- [DH06] D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inf. Theor.*, 47(7):2845–2862, September 2006.
- [GM15] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, pages 829–849, 2015.
- [HK12] Elad Hazan and Satyen Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012, 2012*.
- [HK13] Daniel Hsu and Sham M Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 11–20. ACM, 2013.
- [HSS15] Samuel B. Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 956–1006, 2015.
- [HSSS16] Samuel B. Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 178–191, 2016.

- [JOS05] Nikola Jevtic, Alon Orlitsky, and Narayana P. Santhanam. A lower bound on compression of unknown alphabets. *Theor. Comput. Sci.*, 332(1-3):293–311, 2005.
- [Kle03] Jon M. Kleinberg. An impossibility theorem for clustering. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 463–470. MIT Press, 2003.
- [Las01] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [Las15] Jean Bernard Lasserre. *An introduction to polynomial and semi-algebraic optimization*. Cambridge Texts in Applied Mathematics. Cambridge: Cambridge University Press. , 2015.
- [Lau09] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In Mihai Putinar and Seth Sullivant, editors, *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, pages 157–270. Springer New York, 2009.
- [LLS⁺13] Roi Livni, David Lehavi, Sagi Schein, Hila Nachlieli, Shai Shalev-Shwartz, and Amir Globerson. Vanishing component analysis. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 597–605, 2013.
- [LMSS07] Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.
- [MRT12] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2012.
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *FOCS 2016, to appear*, 2016.
- [MW15] Tengyu Ma and Avi Wigderson. Sum-of-squares lower bounds for sparse PCA. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1612–1620, 2015.
- [Nat95] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.
- [OSZ04] Alon Orlitsky, Narayana P. Santhanam, and Junan Zhang. Universal compression of memoryless sources over unknown alphabets. *IEEE Trans. Information Theory*, 50(7):1469–1481, 2004.
- [Par00] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000.
- [PVMH13] Hristo S Paskov, Robert West, John C Mitchell, and Trevor Hastie. Compressive feature learning. In *Advances in Neural Information Processing Systems*, pages 2931–2939, 2013.
- [Ris78] Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.
- [Sal09] Ruslan Salakhutdinov. *Learning Deep Generative Models*. PhD thesis, University of Toronto, Toronto, Ont., Canada, Canada, 2009. AAINR61080.

- [SS05] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *Learning Theory, 18th Annual Conference on Learning Theory, COLT 2005, Bertinoro, Italy, June 27-30, 2005, Proceedings*, pages 545–560, 2005.
- [SST08] Ohad Shamir, Sivan Sabato, and Naftali Tishby. *Learning and Generalization with the Information Bottleneck*, pages 92–107. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- [TJ89] N. Tomczak-Jaegermann. *Banach-Mazur distances and finite-dimensional operator ideals*. Pitman monographs and surveys in pure and applied mathematics. Longman Scientific & Technical, 1989.
- [TPB00] Naftali Tishby, Fernando C. N. Pereira, and William Bialek. The information bottleneck method. *CoRR*, physics/0004057, 2000.
- [Val84] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [Vap98] Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, 1998.
- [Vu11] Van Vu. Singular vectors under random perturbation. *Random Structures and Algorithms*, 39(4):526–538, 2011.

A Proof of Theorem 2.1

Proof of Theorem 2.1. [MRT12, Theorem 3.1] asserts that with probability at least $1 - \delta$, we have that for every hypothesis $f \in \mathcal{H}$,

$$\text{loss}_{\mathcal{D}}(f) \leq \text{loss}_{\mathcal{S}}(f) + 2\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{\log \frac{1}{\delta}}{2m}}$$

by negating the loss function this gives

$$|\text{loss}_{\mathcal{D}}(f) - \text{loss}_{\mathcal{S}}(f)| \leq 2\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{\log \frac{2}{\delta}}{2m}}$$

and therefore, letting $f^* = \arg \min_{f \in \mathcal{H}} \text{loss}_{\mathcal{D}}(f)$, we have

$$\begin{aligned} \text{loss}_{\mathcal{D}}(\hat{f}_{ERM}) &\leq \text{loss}_{\mathcal{S}}(\hat{f}_{ERM}) + 2\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{\log \frac{1}{\delta}}{2m}} && \text{(by [MRT12, Theorem 3.1])} \\ &\leq \text{loss}_{\mathcal{S}}(f^*) + 2\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{\log \frac{1}{\delta}}{2m}} && \text{(by definition of ERM)} \\ &\leq \text{loss}_{\mathcal{D}}(f^*) + 6\mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{4 \log \frac{1}{\delta}}{2m}} && \text{(using [MRT12, Theorem 3.1] again)} \end{aligned}$$

□

A.1 Low reconstruction error is sufficient for supervised learning

This section observes that low reconstruction error is a sufficient condition for unsupervised learning to allow supervised learning over any future task.

Lemma A.1. Consider any supervised learning problem with respect to distribution \mathcal{D} over $\mathcal{X} \times \mathcal{L}$ that is agnostically PAC-learnable with respect to L -Lipschitz loss function ℓ and with bias γ_1 .

Suppose that unsupervised hypothesis class \mathcal{H} is C -learnable with bias γ_2 over distribution \mathcal{D} and domain \mathcal{X} , by hypothesis $f : \mathcal{X} \mapsto \mathcal{Y}$. Then the distribution $\tilde{\mathcal{D}}_f$ over $\mathcal{Y} \times \mathcal{L}$, which gives the pair $(f(x), y)$ the same measure as \mathcal{D} gives to (x, y) , is agnostically PAC-learnable with bias $\gamma_1 + L\gamma_2$.

Proof. Let $h : \mathcal{X} \mapsto \mathcal{Y}$ be a hypothesis that PAC-learns distribution \mathcal{D} . Consider the hypothesis

$$\tilde{h} : \mathcal{Y} \mapsto \mathcal{L}, \tilde{h}(y) = (h \circ g)(y)$$

Then by definition of reconstruction error and the Lipschitz property of ℓ we have

$$\begin{aligned} \text{loss}_{\tilde{\mathcal{D}}_f}(\tilde{h}) &= \mathbb{E}_{(y,l) \sim \tilde{\mathcal{D}}_f} [\ell(\tilde{h}(y), l)] \\ &= \mathbb{E}_{(y,l) \sim \tilde{\mathcal{D}}_f} [\ell((h \circ g)(y), l)] \\ &= \mathbb{E}_{(x,l) \sim \mathcal{D}} [\ell(h(\tilde{x}), l)] && (\mathcal{D}(x) = \tilde{\mathcal{D}}_f(y)) \\ &= \mathbb{E}_{(x,l) \sim \mathcal{D}} [\ell(h(x), l)] + \mathbb{E}_{(x,l) \sim \mathcal{D}} [\ell(h(\tilde{x}), l) - \ell(h(x), l)] \\ &= \gamma_1 + \mathbb{E}_{(x,l) \sim \mathcal{D}} [\ell(h(\tilde{x}), l) - \ell(h(x), l)] && \text{(PAC learnability)} \end{aligned}$$

$$\begin{aligned}
&\leq \gamma_1 + L \mathbb{E}_{x \sim \mathcal{D}} \|x - \tilde{x}\| && \text{(Lipschitzness of } \ell \circ h) \\
&= \gamma_1 + L \mathbb{E}_{x \sim \mathcal{D}} \|x - g \circ f(x)\| \\
&\leq \gamma_1 + L\gamma_2 && \text{(C-learnability)}
\end{aligned}$$

□

B Proof of Theorem 3.1

Proof of Theorem 3.1. We assume without loss of generality $s = 1$. For $s > 1$ the proof will be identical since one can assume $x^{\otimes s}$ is the data points (and the dimension is raised to d^s).

Let x_1, \dots, x_m be a set of examples $\sim \mathcal{D}^m$. It can be shown that any minimizer of ERM

$$A^* = \arg \min_{A \in \mathbb{R}^{d \times k}} \|x_i - A^\dagger A x_i\|^2 \quad (\text{B.1})$$

satisfies that $(A^*)^\dagger A^*$ is the the projection operator to the subspace of top k eigenvector of $\sum_{i=1}^m x_i x_i^\top$. Therefore ERM (B.1) is efficiently solvable.

According to Theorem 2.1, the ERM hypothesis generalizes with rates governed by the Rademacher complexity of the hypothesis class. Thus, it remains to compute the Rademacher complexity of the hypothesis class for PCA. We assume for simplicity that all vectors in the domain have Euclidean norm bounded by one.

$$\begin{aligned}
\mathcal{R}_S(\mathcal{H}_k^{\text{pca}}) &= \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{(h,g) \in \mathcal{H}_k^{\text{pca}}} \frac{1}{m} \sum_{i \in S} \sigma_i \ell((h,g), x_i) \right] \\
&= \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{A \in \mathbb{R}^{d \times k}} \frac{1}{m} \sum_{i \in S} \sigma_i \|x_i - A^\dagger A x_i\|^2 \right] \\
&= \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{A \in \mathbb{R}^{d \times k}} \frac{1}{m} \sum_{i \in S} \sigma_i \text{Tr} \left((I - A^\dagger A) \left(\sum_{i=1}^m x_i x_i^\top \right) (I - A^\dagger A)^\top \right) \right] \\
&= \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{A \in \mathbb{R}^{d \times k}} \text{Tr} \left((I - A^\dagger A) \left(\frac{1}{m} \sum_{i=1}^m \sigma_i x_i x_i^\top \right) \right) \right].
\end{aligned}$$

Then we apply Holder inequality, and effectively disentangle the part about σ and A :

$$\begin{aligned}
&\mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{A \in \mathbb{R}^{d \times k}} \text{Tr} \left((I - A^\dagger A) \left(\frac{1}{m} \sum_{i=1}^m \sigma_i x_i x_i^\top \right) \right) \right] \\
&\leq \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\sup_{A \in \mathbb{R}^{d \times k}} \|I - A^\dagger A\|_F \left\| \frac{1}{m} \sum_{i=1}^m \sigma_i x_i x_i^\top \right\|_F \right] && \text{(Holder inequality)} \\
&\leq \sqrt{d} \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i x_i x_i^\top \right\|_F \right] && \text{(since } A^\dagger A \text{ is a projection, } \|I - A^\dagger A\| \leq 1.) \\
&\leq \sqrt{d} \mathbb{E}_{\sigma \sim \{\pm 1\}^m} \left[\left\| \frac{1}{m} \sum_{i=1}^m \sigma_i x_i x_i^\top \right\|_F^2 \right]^{1/2} && \text{(Cauchy-Schwarz inequality)} \\
&\leq \sqrt{d} \sqrt{\frac{1}{m^2} \sum_{i=1}^m \langle \sigma_i x_i x_i^\top, \sigma_i x_i x_i^\top \rangle} && \text{(since } \mathbb{E}[\sigma_i \sigma_j] = 0 \text{ for } i \neq j)
\end{aligned}$$

$$\leq \sqrt{d/m}. \quad (\text{by } \|x\| \leq 1)$$

Thus, from Theorem 2.1 we can conclude that the class $\mathcal{H}_k^{\text{pca}}$ is learnable with sample complexity $\tilde{O}(\frac{d}{\varepsilon^2})^6$. \square

B.1 Proof of Theorem 3.2

Theorem 3.2 follows from the following lemmas and the generalization theorem 2.1 straightforwardly.

Lemma B.1. Suppose distribution \mathcal{D} is (k, ε) -regularly spectral decodable. Then for any $\delta > 0$, solving convex optimization (3.3) with non-smooth Frank-Wolfe algorithm [HK12, Theorem 4.4] with $k' = O(\tau^4 k^4 / \delta^4)$ steps gives a solution \hat{R} of rank k' such that $f(\hat{R}) \leq \delta + \varepsilon$.

Lemma B.2. The Rademacher complexity of the class of function $\Phi = \left\{ \|Rx^{\otimes 2} - x^{\otimes 2}\|_{\text{op}} : R \text{ s.t. } \|R\|_{S_1} \leq \tau k \right\}$ with m examples is bounded from above by at most $\mathcal{R}_m(\Phi) \leq 2\tau k \cdot \sqrt{1/m}$

Here Lemma B.2 follows from the fact that $\|Rx^{\otimes 2} - x^{\otimes 2}\|_{\text{op}}$ is bounded above by $2\tau k$ when $\|x\| \leq 1$ and $\|R\|_{S_1} \leq \tau k$. The rest of the section focuses on the proof of Lemma B.1.

Lemma B.1 basically follows from the fact that f is Lipschitz and guarantees of the Frank-Wolfe algorithm.

Proposition B.3. The objective function $f(R)$ is convex and 1-Lipschitz. Concretely, Let $\ell_x(R) = \|Rx^{\otimes 2} - x^{\otimes 2}\|_{\text{op}}$. Then

$$\partial \ell_x \ni (u \otimes v)(x^{\otimes 2})^\top$$

where $\partial \ell_x$ is the set of sub-gradients of ℓ_x with respect to R , and $u, v \in \mathbb{R}^d$ are (one pair of) top left and right singular vectors of $\mathcal{M}(Rx^{\otimes 2} - x^{\otimes 2})$.

Proof. This simply follows from calculating gradient with chain rule. Here we use the fact that $A \in (\mathbb{R}^d)^{\otimes 2}$, the sub-gradient of $\|A\|_{\text{op}}$ contains the vector $a \otimes b$ where a, b are the top left and right singular vectors of $\mathcal{M}(A)$. We can also verify by definition that $(u \otimes v)(x^{\otimes 2})^\top$ is a sub-gradient.

$$\begin{aligned} f(R') - f(R) &\geq (u \otimes v)^\top (R'x^{\otimes 2} - Rx^{\otimes 2}) && (\text{by convexity of } \|\cdot\|_{\text{op}}) \\ &= \langle (u \otimes v)(x^{\otimes 2})^\top, R' - R \rangle. \end{aligned}$$

\square

Now we are ready to prove Lemma B.1.

Proof of Lemma B.1. Since \mathcal{D} is (k, ε) -regularly decodable, we know that there exists a rank- k solution R^* with $f(R^*) \leq \varepsilon$. Since $\|R^*\|_{\text{op}} \leq \tau$, we have that $\|R^*\|_{S_1} \leq \text{rank}(R^*) \cdot \|R^*\|_{\text{op}} \leq \tau k$. Therefore R^* is feasible solution for the objective (3.3) with $f(R^*) \leq \varepsilon$.

By Proposition B.3, we have that $f(R)$ is 1-Lipschitz. Moreover, for any R, S with $\|R\|_{S_1} \leq \tau k, \|S\|_{S_1} \leq \tau k$ we have that $\|R - S\|_F \leq \|R\|_F + \|S\|_F \leq \|R\|_{S_1} + \|S\|_{S_1} \leq 2\tau k$. Therefore the diameter of the constraint set is at most τk .

By [HK12, Theorem 4.4], we have that Frank-Wolfe algorithm returns solution R with $f(R) - f(R^*) \leq \varepsilon + \delta$ in $(\frac{\tau k}{\delta})^4$ iteration. \square

⁶For $\ell > 1$ the sample complexity is $\tilde{O}(d^\ell / \varepsilon^2)$.

C Shorter codes with relaxed objective for Polynomial Spectral Components Analysis

Notations. For a matrix A , let $\sigma_1(A) \geq \sigma_2(A) \geq \dots$ be its singular values. Then the Schatten p -norm, denoted by $\|\cdot\|_{S_p}$, for $p \geq 1$ is defined as $\|A\|_{S_p} = (\sum_i \sigma_i(A)^p)^{1/p}$. For even integer p , an equivalent and simple definition is that $\|A\|_{S_p}^p \triangleq \mathbf{Tr}((A^\top A)^{p/2})$.

In this section we consider the following further relaxation of objective (3.3).

$$\begin{aligned} \text{minimize } f_4(R) &:= \mathbb{E} \left[\|Rx^{\otimes 2} - x^{\otimes 2}\|_{S_p}^2 \right] \\ \text{s.t. } \|R\|_{S_1} &\leq \tau k \end{aligned} \quad (\text{C.1})$$

Since $\|A\|_F \geq \|A\|_{S_4} \geq \|A\|_{S_\infty} = \|A\|$, this is a relaxation of the objective (3.3), and it interpolates between kernel PCA and spectral decoding. Our assumption is weaker than kernel PCA but stronger than spectral decodable.

Definition C.1 (Extension of definition 3.2). We say a data distribution \mathcal{D} is (k, ε) -regularly spectral decodable with $\|\cdot\|_{S_p}$ norm if the error E in equation (5.10) is bounded by $\|E\|_{S_p} \leq \varepsilon$.

We can reduce the length of the code from $O(k^4)$ to $O(k^2)$ for any constant p .

Theorem C.1. Suppose data distribution is (k, ε) -spectral decodable with norm $\|\cdot\|_{S_p}$ for $p = O(1)$, then solving (C.1) using (usual) Frank-Wolfe algorithm gives a solution \hat{R} of $k' = O(k^2\tau^2/\varepsilon^2)$ with $f(R) \leq \varepsilon + \delta$. As a direct consequence, we obtain encoding and decoding pair $(g_A, h_B) \in \mathcal{H}_{k'}^{\text{sa}}$ with $k' = O(k^2\tau^2/\varepsilon^2)$ and reconstruction error $\varepsilon + \delta$.

The main advantage of using relaxed objective is its smoothness. This allows us to optimize over the Schatten 1-norm constraint set much faster using usual Frank-Wolfe algorithm. Therefore the key here is to establish the smoothness of the objective function. Theorem C.1 follows from the following proposition straightforwardly.

Proposition C.2. Objective function f_p (equation (C.1)) is convex and $O(p)$ -smooth.

Proof. Since $\|\cdot\|_{S_p}$ is convex and composition of convex function with linear function gives convex function. Therefore, $\|Rx^{\otimes 2} - x^{\otimes 2}\|_{S_p}$ is a convex function. The square of a non-negative convex function is also convex, and therefore we proved that f_p is convex. We prove the smoothness by first showing that $\|A\|_{S_p}^2$ is a smooth function with respect to A . We use the definition $\|A\|_{S_p}^p = \mathbf{Tr}((A^\top A)^{p/2})$. Let E be a small matrix that goes to 0, we have

$$\|A + E\|_{S_p}^p = \mathbf{Tr}((A^\top A)^{p/2}) + T_1 + T_2 + o(\|E\|_F^2) \quad (\text{C.2})$$

where T_1, T_2 denote the first order term and second order term respectively. Let $U = A^\top E + E^\top A$ and $V = A^\top A$. We note that T_2 is a sum of the traces of matrices like $UVUVU^{p/2-2}$. By Lieb-Thirring inequality, we have that all these term can be bounded by $\mathbf{Tr}(U^{p/2-2}V^2) = 2\mathbf{Tr}((A^\top A)^{p/2-2}A^\top EE^\top A) + 2\mathbf{Tr}((A^\top A)^{p/2-2}A^\top EA^\top E^\top)$. For the first term, we have that

$$\mathbf{Tr}((A^\top A)^{p/2-2}A^\top EE^\top A) \leq \|(AA^\top)^{(p-2)/4}E\|^2 \leq \|(AA^\top)^{(p-2)/4}\|_{S_\infty}^2 \|E\|_F^2 = \|A\|_{S_\infty}^{p-2} \|E\|_F^2$$

where in the first inequality we use Cauchy-Schwarz. Then for the second term we have

$$\begin{aligned} \mathbf{Tr}((A^\top A)^{p/2-2}A^\top EA^\top E^\top) &\leq \|(A^\top A)^{(p-2)/4}E\|_F \|AE(A^\top A)^{(p-4)/4}\|_F \\ &\leq \|(A^\top A)^{(p-2)/4}E\|_F^2 \quad (\text{by Lieb-Thirring inequality}) \\ &\leq \|(AA^\top)^{(p-2)/4}\|^2 \|E\|_F^2 = \|A\|_{S_\infty}^{p-2} \|E\|_F^2 \end{aligned} \quad (\text{C.3})$$

Therefore, finally we got

$$T_2 \leq O(p^2) \cdot \|A\|_{S_{p-2}}^{p-2} \|E\|_F^2 \quad (\text{C.4})$$

Therefore, we complete the proof by having,

$$\begin{aligned} \|A + E\|_{S_p}^2 &\leq (\|A\|_{S_p}^p + T_1 + T_2 + o(\|E\|^2))^{2/p} \leq \|A\|_{S_p} (1 + T_1' + \frac{2}{p\|A\|_{S_p}^{p/2}} T_2) + o(\|E\|^2) \\ &\leq \|A\|_{S_p}^2 + T_1'' + O(p)\|E\|^2 + o(\|E\|^2) \end{aligned} \quad \begin{array}{l} \text{(by } (1+x)^{p/2} \leq 1 + 2x/p + o(\|x\|^2)\text{)} \\ \text{(by equation (C.3))} \end{array}$$

□

D Toolbox

Lemma D.1. Let $p \geq 2$ be a power of 2 and $u = [u_1, \dots, u_n]$ and $v = [v_1, \dots, v_n]$ be indeterminants. Then there exists SoS proof,

$$\vdash \left(\sum_j u_j v_j^{p-1} \right)^p \leq \left(\sum_i u_i^p \right) \left(\sum v_i^p \right)^{p-1} \quad (\text{D.1})$$

Proof Sketch. The inequality follows from repeated application of Cauchy-Schwarz. For example, for $p = 4$ we have

$$\begin{aligned} \vdash \left(\sum_i u_i^4 \right) \left(\sum v_i^4 \right)^3 &\geq \left(\sum_i u_i^2 v_i^2 \right)^2 \left(\sum v_i^4 \right)^2 && \text{(by Cauchy-Schwarz)} \\ &\geq \left(\sum_i u_i v_i^3 \right)^4 && \text{(by Cauchy-Schwarz again)} \end{aligned}$$

For $p = 2^s$ with $s > 2$, the statement can be proved inductively. □