Quantum Entanglement, Sum of Squares, and the Log Rank Conjecture*

Boaz Barak
Harvard University
Cambridge, MA, USA
kothari@cs.princeton.edu

Pravesh K. Kothari
Princeton University and the IAS
Princeton, NJ, USA
b@boazbarak.org

David Steurer
Cornell University and IAS
Ithaca, NY, USA
dsteurer@cs.cornell.edu

ABSTRACT

For every constant $\varepsilon > 0$, we give an $\exp(\tilde{O}(\sqrt{n}))$ -time algorithm for the 1 vs $1 - \varepsilon$ *Best Separable State (BSS)* problem of distinguishing, given an $n^2 \times n^2$ matrix M corresponding to a quantum measurement, between the case that there is a separable (i.e., non-entangled) state ρ that M accepts with probability 1 , and the case that every separable state is accepted with probability at most $1 - \varepsilon$. Equivalently, our algorithm takes the description of a subspace $\mathcal{W} \subseteq \mathbb{F}^{n^2}$ (where $\mathbb{F}$ can be either the real or complex field) and distinguishes between the case that $\mathcal{W}$ contains a rank one matrix, and the case that every rank one matrix is at least ε far (in ℓ_2 distance) from $\mathcal{W}$.

To the best of our knowledge, this is the first improvement over the brute-force $\exp(n)$ -time algorithm for this problem. Our algorithm is based on the *sum-of-squares* hierarchy and its analysis is inspired by Lovett's proof (STOC '14, JACM '16) that the communication complexity of every rank- n Boolean matrix is bounded by $\tilde{O}(\sqrt{n})$.

CCS CONCEPTS

• **Theory of computation** → Computational complexity and cryptography; Design and analysis of algorithms; Approximation algorithms analysis; Continuous optimization; Semidefinite programming; Data structures design and analysis;

KEYWORDS

Best Separable State Problem, Sum-of-Squares Semidefinite Programming Hierarchy, Polynomial Reweighting

ACM Reference format:

Boaz Barak, Pravesh K. Kothari, and David Steurer. 2017. Quantum Entanglement, Sum of Squares, and the Log Rank Conjecture. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC'17)*, 14 pages.

*Full Version here: <https://arxiv.org/abs/1701.06321>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
STOC'17, Montreal, Canada

© 2017 ACM. 978-1-4503-4528-6/17/06...\$15.00
DOI: 10.1145/3055399.3055488

DOI: 10.1145/3055399.3055488

1 INTRODUCTION

Entanglement is one of the more mysterious and subtle phenomena of quantum mechanics. The formal definition is below (Definition 1.2), but roughly speaking, a quantum state ρ on two systems A and B is *entangled* if a quantum measurement of one system can affect the other system. A non-entangled state is called *separable*. This type of "spooky interaction at a distance" is responsible for many of the more counter-intuitive features of quantum mechanics. Entanglement is also used by all algorithms for quantum computers that obtain speedups over the best known classical algorithms, and it may be necessary for such speedups [Vid03].

One of the ways in which the complexity of entanglement is manifested is that even given the full description of a quantum state ρ as a density matrix, there is no known efficient algorithm for determining whether ρ is entangled or not. Indeed, the best known algorithms take time which is *exponential* in the dimension of the state (which itself is exponential in the number of underlying qubits). This is in contrast to the classical case, where there is an efficient algorithm that given a probability distribution μ over a universe $A \times B$, can check whether or not μ is a *product distribution* by simply computing the rank of μ when viewed as a matrix.

Given the inherently probabilistic and noisy setting of quantum computing, arguably the right question is not to determine entanglement exactly, but rather to distinguish between the case that a state ρ is separable, and the case that it is ε -far from being separable, in the sense that there exists some *measurement* M that accepts ρ with probability p but accepts every separable state with probability at most $p - \varepsilon$. This problem is known as the *Quantum Separability Problem* with parameter ε . Gharibian [Gha10], improving on Gurvits [Gur03], showed that this problem is NP hard when ε is inversely polynomial in the dimension of the state. Harrow and Montanaro [HM13] showed that, assuming the Exponential Time Hypothesis, there is no $n^{o(\log n)}$ time algorithm for this problem for ε which is a small constant.

A tightly related problem, which is the one we focus on in this paper, is the *Best Separable State (BSS)* problem.¹ In the BSS problem

¹Using the connection between optimization and separation oracles in convex programming, one can convert a sufficiently good algorithm for the search variant of one of these problems to the other. See [HM13, Sec. 4.2] for a thorough discussion of the relations between these and many other problems.

the input is a measurement M on a two part system and two numbers $1 \geq c > s \geq 0$ and the goal is to distinguish between the YES case that there is a separable state that M accepts with probability at least c and the NO case that M accepts every separable state with probability at most s . In particular, certifying that a particular measurement M satisfies the NO case is extremely useful since it implies that M can serve as *entanglement witness* [HHH96, LKCH00], in the sense that achieving acceptance probability with M larger than s certifies the presence of entanglement in a state. Such entanglement witnesses are used to certify entanglement in experiments and systems such as candidate computing devices [Ved08], and so having an efficient way to certify that they are sound (do not accept separable states) can be extremely useful.

Similarly to the quantum separability problem, the BSS problem is NP hard when $c - s = 1/\text{poly}(n)$ [BT09, Gur03] and Harrow and Montanaro [HM13, Corollary 13(i)] show that (assuming the ETH) there is no $n^{o(\log n)}$ time algorithm for $\text{BSS}_{1,1/2}$. An outstanding open question is whether the [HM13] result is *tight*: whether there is a quasi-polynomial time algorithm for $\text{BSS}_{c,s}$ for some constants $1 \geq c > s \geq 0$. This question also has a complexity interpretation. A measurement on a two part system can be thought of as a *verifier* (with hardwired input) that interacts with two provers. Requiring the state to be *separable* corresponds to stipulating that the two provers are not entangled. Thus it is not hard to see that an algorithm for $\text{BSS}_{c,s}$ corresponds to an algorithm for deciding all languages in the complexity class $\text{QMA}(2)$ of *two prover quantum Merlin Arthur* systems with corresponding completeness and soundness parameters c and s respectively. In particular a quasi-polynomial time algorithm for $\text{BSS}_{0.99,0.5}$ would imply that $\text{QMA}(2) \subseteq \text{EXP}$, resolving a longstanding problem in quantum complexity.²

In 2004, Doherty, Parrilo and Spedalieri [DPS04] proposed an algorithm for the BSS problem based on the *Sum of Squares* semidefinite programming hierarchy [Par00, Las01]. It is not known whether this algorithm can solve the $\text{BSS}_{c,s}$ problem (for constants $c > s$) in quasi-polynomial time. However Brandão, Christandl and Yard [BaCY11] showed that it runs in quasi-polynomial time when the measurement M is restricted to a special class of measurements known as *one-way local operations and classical communications* (1-LOCC). Brandão and Harrow [BH15] showed that similar performance for these types of measurements can be achieved by an algorithm based on searching on an appropriately defined ε -net.

1.1 Non Quantum Motivations

The BSS problem is actually quite natural and well motivated from classical considerations. As we'll see in Section 2 below, it turns out that at its core lies the following problem:

Definition 1.1 (Rank one vector in subspace problem). Let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ and $\varepsilon > 0$. The ε rank one vector problem over $\mathbb{F}$ is the task

²For more on information on this problem and its importance, see the presentations in the recent workshop <http://qma2016.quics.umd.edu/> that was dedicated to it.

of distinguishing, given a linear subspace $\mathcal{W} \subseteq \mathbb{F}^{n^2}$, between the case that there is a nonzero rank one matrix $L \in \mathcal{W}$ and the case that $\|L - M\|_F \geq \varepsilon \|L\|_F$ for every rank one L and $M \in \mathcal{W}$.³

This is arguably a natural problem in its own right. While solving this problem exactly (i.e., determining if there is a rank one solution to a set of linear equations) is the same as the NP hard task of solving *quadratic equations*, it turns out that we can obtain non-trivial algorithmic results by considering the above notion of approximation. Indeed, our main result implies an $\exp(\tilde{O}(\sqrt{n}))$ time algorithm for this problem for any constant $\varepsilon > 0$ in both the real and complex cases.

1.2 Our Results

In this work we give a $2^{\tilde{O}(\sqrt{n})}$ time algorithm for the $\text{BSS}_{1,s}$ problem for every constant $s < 1$. We now make the necessary definitions and state our main result.⁴

Definition 1.2. A *quantum state* on a system of m elementary states (e.g., a log m -qubit register) is an $m \times m$ complex Hermitian matrix ρ (known as a *density matrix*) such that $\text{Tr } \rho = 1$. A quantum state ρ is *pure* if it is of the form $\rho = ww^*$ for some unit vector $w \in \mathbb{C}^m$. Otherwise we say that ρ is *mixed*. Note that every mixed state ρ is a convex combination of pure states.

If $m = n^2$, and we identify $[m]$ with $[n] \times [n]$ then an m -dimension pure quantum state $\rho = ww^* \in \mathbb{C}^{n^2}$ is *separable* if the vector $w \in \mathbb{C}^m$ is equal to uv^* for some $u, v \in \mathbb{C}^n$. A general state ρ is *separable* if it is a convex combination of separable pure states. That is, $\rho = \mathbb{E}(uv^*)(uv^*)^*$ where the expectation is taken over a distribution supported over pairs of unit vectors $u, v \in \mathbb{C}^n$. A state that is not separable is called *entangled*.

A quantum *measurement operator* is an $m \times m$ complex Hermitian matrix M such that $0 \leq M \leq I$. The probability that a measurement M accepts a state ρ is $\text{Tr}(\rho M)$.

THEOREM 1.3 (MAIN RESULT). *For every $s < 1$, there is a $2^{\tilde{O}(\sqrt{n})}$ time algorithm, based on $\tilde{O}(\sqrt{n})$ rounds of the sos hierarchy, that on input an $n^2 \times n^2$ measurement operator M , distinguishes between the following two cases:*

- YES: There exists a separable state $\rho \in \mathbb{C}^{n^2 \times n^2}$ such that $\text{Tr}(\rho M) = 1$.
- NO: For every separable $\rho \in \mathbb{C}^{n^2 \times n^2}$, $\text{Tr}(\rho M) \leq s$

To our knowledge, this algorithm is the first for this problem that beats the brute force bound of $2^{O(n)}$ time for general measurements.

³For a $k \times m$ matrix A , we denote by $\|A\|_F$ its *Frobenius* norm, defined as $\sqrt{\sum_{i,j} |A_{i,j}|^2} = \text{Tr}(AA^*)^{1/2}$, which is the same as taking the ℓ_2 norm of the matrix when considered as an km -dimensional vector.

⁴For the sake of accessibility, as well as to emphasize the connections with non-quantum questions, we use standard linear algebra notation in this paper as opposed to Dirac's ket notation that is more common in quantum mechanics. A vector u is a column vector unless stated otherwise, and u^* denotes the complex conjugate transpose of the vector u . If u is real, then we denote its transpose by u^T . See the lecture notes [BCHW16] for a more complete coverage of separability and entanglement.

Like the algorithms of [DPS04, BaCY11], our algorithm is based on the *sum of squares* SDP hierarchy, but we introduce new techniques for analyzing it that we believe are of independent interest. As we discuss in Section 8, it is a fascinating open question to explore whether our techniques can be quantitatively strengthened to yield faster algorithms and/or extended for other problems such as the 2 to 4 norm and small set expansion problems, that have been shown to be related to the BSS problem by [BBH⁺12] (albeit in a different regime of parameters than the one we deal with in this work). As we remark below, this question seems related to other longstanding open questions in computer science and in particular to the *log rank conjecture* in communication complexity [LS88].

Remark 1.4 (Imperfect completeness). We state our results for the case of perfect completeness for simplicity, but all of the proofs extend to the case of “near perfect completeness” where in the YES case we replace the condition $\text{Tr}(\rho\mathcal{M}) = 1$ with the condition $\text{Tr}(\rho\mathcal{M}) = 1 - \frac{1}{n}$ (see the proof of Theorem 4.2). It is an interesting open problem to find out whether our results can extend to the setting where in the YES case $\text{Tr}(\rho\mathcal{M}) = 1 - \varepsilon$ for some absolute constant ε . We conjecture that this is indeed the case.

Remark 1.5 (Real vs complex numbers). While the natural setting for quantum information theory is the *complex numbers*, much of the power and interest already arises in the case of the real numbers, which is more natural for the sos algorithm (though it does have complex-valued generalization). For our purposes, there’s no difference between the real and the complex cases - we give a reduction from the complex case to the real case in Section B of the Appendix. Thus, from now on, we will focus solely on the case that all operators, subspaces, matrices are *real*.

2 OUR TECHNIQUES

Our algorithm follows a recent paradigm of constructing rounding algorithms for the sum of squares sdp by considering its solutions as “pseudo distributions” [BKS16]. These can be thought of as capturing the uncertainty that a computationally bounded solver has about the optimal solution of the given problem, analogously to the way that probability distributions model uncertainty in the classical information-theoretic Bayesian setting.

Somewhat surprisingly, our main tool in analyzing the algorithm are techniques that arose in proof of the currently best known upper bound for the *log rank conjecture* [LS88]. This conjecture has several equivalent formulations, one of which is that every $N \times N$ matrix A with Boolean (i.e., 0/1) entries and rank at most n , contains a submatrix of size at least $2^{-\text{poly log}(n)} N \times 2^{-\text{poly log}(n)} N$ that is of rank one.⁵ The best known bound on the log rank conjecture is

⁵The original formulation of the log rank conjecture is that every such matrix has communication complexity at most $\text{poly log}(n)$, and Nisan and Wigderson [NW94] showed that this is equivalent to the condition that such matrices contains a monochromatic submatrix of the above size. Every monochromatic submatrix is rank one, and every rank one submatrix of size $s \times s$ of a Boolean valued matrix contains a monochromatic submatrix of size at least $\frac{s}{2} \times \frac{s}{2}$.

by Lovett [Lov14] who proved that every such matrix contains a submatrix of size at least $2^{-\tilde{O}(\sqrt{n})} N \times 2^{-\tilde{O}(\sqrt{n})} N$.

Our algorithm works by combining the following observations:

- (1) Lovett’s proof can be generalized to show that *every* $N \times N$ rank n real (or complex) matrix A (not necessarily with Boolean entries) contains a $2^{-\tilde{O}(\sqrt{n})} N \times 2^{-\tilde{O}(\sqrt{n})} N$ submatrix that is *close* to rank one in Frobenius norm.
- (2) If μ is an *actual* distribution over solutions to the sos program for the BSS problem on dimension n , then we can transform μ into an $N \times N$ rank n matrix $A = A(\mu)$ such that extracting an approximate solution from A in time $2^{\tilde{O}(k)}$ can be done if A contains an approximately rank one submatrix of size at least $2^{-k} N \times 2^{-k} N$.
- (3) Moreover all the arguments used to establish steps 1 and 2 above can be encapsulated in the sum of squares framework, and hence yield an algorithm that extracts an approximately optimal solution to the BSS problem from a degree $\tilde{O}(\sqrt{n})$ pseudo-distribution μ that “pretends” to be supported over exact solutions.

Thus, even though in the sos setting there is no actual distribution μ , and hence no actual matrix A , we can still use structural results on this “fake” (or “pseudo”) matrix A to obtain an *actual* rounding algorithm. We view this as a demonstration of the power of the “pseudo distribution” paradigm to help in the discovery of new algorithms, that might not seem as natural without placing them in this framework.

2.1 Rounding from Rank-One Reweightings

We now give a more detailed (yet still quite informal) overview of the proof. As mentioned above, we focus on the case that the $n^2 \times n^2$ measurement matrix $\mathcal{M}$ is *real* (as opposed to *complex*) valued.

Let $\mathcal{W} \subseteq \mathbb{R}^{n^2}$ be the subspace of vectors X such that $X^\top \mathcal{M} X = \|X\|^2$ (this is a subspace since $\mathcal{M} \leq I$ and hence $\mathcal{W}$ is the eigenspace of $\mathcal{M}$ corresponding to the eigenvalue 1). We pretend that the sos algorithm yields a distribution μ over rank one matrices of the form $X = uv^\top$ such that $X \in \mathcal{W}$. When designing a rounding algorithm, we only have access to *marginals* of μ , of the form $\mathbb{E}_\mu f(X)$ for some “simple” function f (e.g., a low degree polynomial). We need to show that we can use such “simple marginals” of μ to extract a single rank one matrix $u_0 v_0^\top$ that has large projection into $\mathcal{W}$.

We start with the following simple observation:

LEMMA 2.1. *If μ is a distribution over matrices X in a subspace $\mathcal{W} \subseteq \mathbb{R}^{n^2}$ such that the expectation $\mathbb{E}_\mu X$ is approximately rank one, in the sense that $\|L - \mathbb{E}_\mu X\|_F \leq \varepsilon \|L\|_F$ for some rank one matrix L , then $\text{Tr}(\mathcal{M}\rho) \geq 1 - 2\varepsilon^2$ where ρ is the pure separable state $\rho = LL^\top / \|L\|_F^2$.*

PROOF. Since μ is supported over matrices in $\mathcal{W}$, $\mathbb{E}_\mu X$ is in $\mathcal{W}$. But this means that the ℓ_2 (i.e., Frobenius) norm distance of L to the subspace $\mathcal{W}$ is at most $\varepsilon \|L\|_F$. Since $\text{Tr}(XX^\top \mathcal{M}) = \text{Tr}(X^\top \mathcal{M} X) =$

$\|X\|_F^2$ for every $X \in \mathcal{W}$, the value $\text{Tr}(LL^\top M)$ will be at least as large as the norm squared of the projection of L to $\mathcal{W}$. $\square$

In particular this means that if we were lucky and the condition of Lemma 2.1's statement occurs, then it would be trivial for us to extract from the expectation $\mathbb{E}_\mu X$ (which is a very simple marginal) a rank one matrix that is close to $\mathcal{W}$, and hence achieves probability $1 - \varepsilon$ in the measurement $\mathcal{M}$. Note that even if every matrix in the support of μ has unit norm, the matrix L could be of significantly smaller norm. We just need that there is some dimension-one subspace on which the cancellations among these matrices are significantly smaller than the cancellations in the rest of the dimensions.

Of course there is no reason we should be so lucky, but one power that the marginals give us is the ability to *reweight* the original distribution μ . In particular, for every "simple" non-negative function $\zeta : \mathbb{R}^{n^2} \rightarrow \mathbb{R}_+$, we can compute the marginal $\mathbb{E}_{\mu_\zeta} X$ where μ_ζ is the distribution over matrices where $\mathbb{P}_{\mu_\zeta}[X]$ (or $\mu_\zeta(X)$ for short) is proportional to $\zeta(X)\mu(X)$. A priori in the degree k sos algorithm we are only able to reweight using functions ζ that are polynomials of degree at most k , but for the purposes of this overview, let us pretend that we can reweight using any function that is not too "spiky" and make the following definition:

Definition 2.2. Let μ be a probability distribution. We say that a probability distribution μ' is a *k-deficient reweighting* of μ if $\Delta_{KL}(\mu' \parallel \mu) \leq k$ where $\Delta_{KL}(\mu' \parallel \mu)$ denotes the Kullback-Leibler divergence of μ' and μ , defined as $\mathbb{E}_{X \sim \mu'} \log(\mu'(X)/\mu(X))$.

Thus at least on a "moral level", the following theorem should be helpful for proving our main result:

THEOREM 2.3 (RANK ONE REWEIGHTING). *Let μ be any distribution over rank one $n \times n$ matrices and $\varepsilon > 0$. Then there exists an $\sqrt{n} \text{poly}(1/\varepsilon)$ -deficient reweighting μ' of μ and a rank one matrix L such that*

$$\|L - \mathbb{E}_{\mu'} X\|_F \leq \varepsilon \|L\|_F$$

One of the results of this paper is a proof of Theorem 2.3 (see Section 2.3). It turns out that this can be done using ideas from the works on the log rank conjecture.

2.2 From Monochromatic Rectangles to Rank-one Reweightings

What does Theorem 2.3 has to do with the log rank conjecture? To see the connection let us imagine that the distribution μ is *flat* in the sense that it is a uniform distribution over rank one matrices $\{u_1 v_1^\top, \dots, u_N v_N^\top\}$ (this turns out to be essentially without loss of generality) and consider the $n \times N$ matrices U and V whose columns are $u_1, \dots, u_N$ and $v_1, \dots, v_N$ respectively. The $n \times n$ matrix $\mathbb{E}_\mu u_i v_i^\top$ is proportional to $UV^\top$. This matrix has the same spectrum (i.e., singular values) as the $N \times N$ matrix $U^\top V$. Hence, $UV^\top$ is close to a rank one matrix if and only if $U^\top V$ is, since in both cases this happens when the square of the top singular value

dominates the sum of the squares of the rest of the singular values. Now a flat distribution μ' with $\Delta_{KL}(\mu' \parallel \mu) \leq k$ corresponds to the uniform distribution over $\{u_i v_i^\top\}_{i \in I}$ where $I \subseteq [N]$ satisfies $|I| \geq 2^{-k} N$. We can see that $\mathbb{E}_{\mu'} u_i v_i^\top$ will be approximately rank one if and only if the submatrix of $U^\top V$ corresponding to I is approximately rank one. Using these ideas it can be shown that Theorem 2.3 is equivalent to the following theorem:⁶

THEOREM 2.4 (RANK ONE REWEIGHTING—DUAL FORMULATION). *Let A be any $N \times N$ matrix of rank at most n . Then there exists a subset $I \subseteq [N]$ with $|I| \geq \exp(-\sqrt{n} \text{poly}(1/\varepsilon))N$ and a rank one matrix L such that*

$$\|L - A_{I,I}\|_F \leq \varepsilon \|L\|_F$$

where $A_{I,I}$ is the submatrix corresponding to restricting the rows and columns of A to the set I .

One can think of Theorem 2.4 as an approximate and robust version of Lovett's result [Lov14] mentioned above. Lovett showed that every $N \times N$ matrix of rank n with *Boolean* entries has a $2^{-\tilde{O}(\sqrt{n})} N \times 2^{-\tilde{O}(\sqrt{n})} N$ submatrix that is of exactly rank 1. We show that the condition of Booleanity is not needed if one is willing to relax the conclusion to having a submatrix that is only *approximately* rank 1. It is of course extremely interesting in both cases whether the bound of $\tilde{O}(\sqrt{n})$ can be improved further, ideally all the way to $\text{polylog}(n)$. In the Boolean setting, such a bound might prove the log rank conjecture,⁷ while in our setting such a bound (assuming it extends to "pseudo matrices") would yield a quasipolynomial time algorithm for BSS, hence showing that $\text{QMA}(2) \subseteq \text{EXP}$. It can be shown that as stated, Theorem 2.3 is tight. However there are different notions of being "close to rank one" that could be useful in both the log-rank and the quantum separability setting, for which there is hope to obtain substantially improved quantitative bounds. We discuss some of these conjectural directions in Section 8.

2.3 Overview of the Proof

In the rest of this technical overview, we give a proof sketch of Theorem 2.4 and then discuss how the proof can be "lifted" to hold in the setting of sum of square pseudo-distributions. The condition that a matrix A is of rank n is the same as that $A = UV^\top$ where U, V are two $n \times N$ matrices with columns $u_1, \dots, u_N$ and $v_1, \dots, v_N$ respectively (i.e., $A_{i,j} = \langle u_i, v_j \rangle$ for all $i, j \in [N]$). We will restrict our attention to the case that all the columns of U and V are of unit norm. (This restriction is easy to lift and anyway holds automatically in our intended application.) In this informal overview, we

⁶To show this formally we use the fact that by Markov, every distribution μ' with $\Delta_{KL}(\mu' \parallel U_{[N]}) = \log N - H(\mu') = k$ is ε -close to a distribution with min entropy $\log N - O(k/\varepsilon)$ and every distribution of the latter type is a convex combination of flat distributions of support at least $N 2^{-O(k/\varepsilon)}$.

⁷We note a caveat that this depends on the notion of "approximate" used. Gavinsky and Lovett [GL14] showed that to prove the log rank conjecture it suffices to find a rank n Boolean matrix a rectangle of measure $\exp(-\text{polylog}(n))$ that is *nearly monochromatic* in the sense of having a $1 - 1/O(n)$ fraction of its entries equal. In this paper we are more concerned with rectangles whose distance to being rank one (or monochromatic) is some $\varepsilon > 0$ that is only a small constant or $1/\text{polylog}(n)$.

also restrict attention to the *symmetric* case, in which $A = A^\top$ and can be written as $A = UU^\top$ and also assume that U is *isotropic*, in the sense that $\mathbb{E}_{i \in [N]} u_i u_i^\top = \frac{1}{n} \text{Id}$.

Our inspiration is Lovett's result [Lov14] which establishes a stronger conclusion for Boolean matrices. In particular, our proof follows Rothvoß's proof [Rot14] of Lovett's theorem, though the non-Boolean setting does generate some non-trivial complications. The $N \times N$ matrix A satisfies that $A_{i,j} = \langle u_i, u_j \rangle$. An equivalent way to phrase our goal is that we want to find a subset $I \subseteq [N]$ over the indices such that:

- (i) $|I| \geq \exp(-\tilde{O}(\sqrt{n}))N$.
- (ii) If $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ are the eigenvalues of $\mathbb{E}_{i \in I} u_i u_i^\top$ then $\varepsilon^2 \lambda_1^2 \geq \sum_{j=2}^n \lambda_j^2$

We will choose the set I *probabilistically* and show that (i) and (ii) above hold in *expectation*. It is not hard to use standard concentration of measure bounds to then deduce the desired result but we omit these calculations from this informal overview.

Our initial attempt for the choice of I is simple, and is directly inspired by [Rot14]. We choose a random standard Gaussian vector $g \in N(0, \frac{1}{n} \text{Id})$ (i.e., for every i , g_i is an independent standard Gaussian of mean zero and variance $1/n$). We then define $I_g = \{i : \langle g, u_i \rangle \geq \sqrt{k/n}\}$ where $k = \tilde{O}(\sqrt{n})$ is a parameter to be chosen later. Since u_i is a unit vector, $\langle g, u_i \rangle$ is a Gaussian of variance $1/n$, and so for every i , the probability that $i \in I_g$ is $\exp(-O(k))$ hence satisfying (i) in expectation.

The value λ_1 of $\mathbb{E}_{i \in I} u_i u_i^\top$ will be at least $\Omega(k/n)$ in expectation. Indeed, we can see that the Gaussian vector g that we choose (which satisfies $\|g\|^2 = 1 \pm o(1)$ with very high probability) will satisfy that $g^\top (\mathbb{E}_{i \in I_g} u_i u_i^\top) g = \mathbb{E}_{i \in I_g} \langle u_i, g \rangle^2 \geq k/n$ and hence in expectation the top eigenvalue of $\mathbb{E}_{i \in I_g} u_i u_i^\top$ will be at least $(1 - o(1))k/n$.

So, if we could only argue that in expectation it will hold that $\sum_{j=1}^n \lambda_j^2 \ll k^2/n^2 = \text{polylog}(n)/n$ then we'd be done. Alas, this is not necessarily the case. However, if this does fail, we can see that we have made progress, in the sense that by restricting to the indices in I we raised the Frobenius norm of $\mathbb{E}_{i \in I} u_i u_i^\top$ from the previous value of $1/n$ (under the assumption that U was isotropic) to $\text{polylog}(n)/n$. Our idea is to show that this holds in general: we can select a Gaussian vector g and define the set I_g as above such that by restricting to the indices in I_g we either get an approx rank one matrix or we increase the Frobenius norm of our expectation matrix by at least an $(1 + \varepsilon)$ factor for an appropriately chosen $\varepsilon > 0$. Since the latter cannot happen more than $\log n/\varepsilon$ times, the final set of indices still has measure $\exp(-\tilde{O}(\sqrt{n}))$.

In further rounds, if our current set of indices is I and the matrix (after subtracting from each vector u_i its expectation) $U_I = \mathbb{E}_{i \in I} u_i u_i^\top = \sum_{j=1}^n \lambda_j v_j v_j^\top$ is not approximately rank one, then rather than choosing g as a standard Gaussian, we choose it from the distribution $N(0, U_I)$ where we use U_I as the covariance matrix. The expected norm of g is simply $\text{Tr}(U_I)$ which equals 1. For every i , the random variable $\langle u_i, g \rangle$ is a Gaussian with mean zero and variance $\sum_{j=1}^n \langle u_i, v_j \rangle \lambda_j$. But for every j in expectation

over i , $\mathbb{E} \langle u_i, v_j \rangle^2 = \lambda_j$ and so it turns out that we can assume that this random variable has variance $\sum \lambda_j^2 = \|U_I\|_F^2$.

This means that if we choose $I' = \{i \in I : \langle u_i, g \rangle \geq \sqrt{k} \|U_I\|_F\}$ we get a subset of I with measure $\exp(-O(k))$. But now the new matrix $U_{I'} = \mathbb{E}_{i \in I'} u_i u_i^\top$ will have an eigenvalue of at least $k \|U_I\|_F^2$ magnitude which is much larger than $\|U_I\|_F$ since we chose $k \gg \sqrt{n}$. Hence $U_{I'}$ has significantly larger Frobenius norm than U_I . The above arguments can be made precise and yield a proof of Theorem 2.4 and thus also Theorem 2.3.

2.4 Rectangle Lemma for Pseudo-distributions

The above is sufficient to show that given $N \times n$ matrices $U = (u_1 | \dots | u_N)$ and $V = (v_1 | \dots | v_n)$ (which we view as inducing a distribution over rank one matrices by taking $u_i v_i^\top$ for a random i), we can condition on a not too unlikely event (of probability $\exp(-\tilde{O}(\sqrt{n}))$) to obtain that $\mathbb{E} u_i v_i^\top$ is roughly rank one. But in the sos setting we are *not* given such matrices. Rather we have access to an object called a "pseudo-distribution" μ which we behaves to a certain extent as if it is such a distribution, but for which it is not actually the case. In particular, we are not able to sample from μ , or condition it on arbitrary events, but rather only compute $\mathbb{E}_\mu f(X)$ for polynomials f of degree at most $\tilde{O}(\sqrt{n})$, and even these expectations are only "pseudo expectations" in the sense that they do not need to correspond to any actual probability distribution.

To lift the arguments above to the sos setting, we need to first show that if μ was an actual distribution, then we could perform all of the above operations using only access to $\tilde{O}(\sqrt{n})$ degree moments of μ . Then we need to show that our *analysis* can be captured by the degree $\tilde{O}(\sqrt{n})$ sos proof systems. Both these steps, which are carried out in Sections 6 and 7 of this paper, are rather technical and non-trivial, and we do not describe them in this overview.

For starters, we need to move from *conditioning* a probability distribution to *reweighting* it. All of our conditioning procedures above had the form of restricting to i 's such that $\xi(i) \geq \sqrt{k}$ where $\xi(i)$ was probabilistically chosen so that for every i $\xi(i)$ is a mean zero and standard deviation one random variable satisfying $\mathbb{P}[\xi(i) = \ell] = \exp(-\Theta(\ell^2))$. We replace this conditioning by *reweighting* the distribution i with the function $\zeta(i) = \exp(\sqrt{k} \xi(i))$. Note that iterative conditioning based on functions $\xi_1, \dots, \xi_t$ can be replaced with reweighting by the product function $\zeta_1, \dots, \zeta_t$. We then show that these ζ_j functions can be approximated by polynomials of $\tilde{O}(k)$ degree.

The arguments above allow us to construct a rounding algorithm that at least makes sense syntactically, in the sense that it takes the $\tilde{O}(\sqrt{n})$ degrees moments of μ and produces a rank one matrix that is a candidate solution to the original matrix. To analyze this algorithm, we need to go carefully over our analysis before, and see that all the arguments used can be embedded in the sos proof system with relatively low degree. Luckily we can rely on the recent body of works that establishes a growing toolkit of techniques to show such embeddings [BKS16].

3 PRELIMINARIES

We use the standard $O(\cdot)$ and $\Omega(\cdot)$ notation to hide absolute multiplicative constants. We define $\mathbb{S}^{n-1}$ to be the $n-1$ dimensional unit sphere $\{x \in \mathbb{R}^n : \sum_i |x_i|^2 = 1\}$. We use the following definitions related the sum of squares (sos) algorithm; see [BKS16] for a more in depth treatment.

Definition 3.1. Let $n \in \mathbb{N}$ and $[x_1, x_2, \dots, x_n]_d$ be the subspace of all n -variate real polynomials of degree at most d . A *degree- d pseudo distribution* μ over $\mathbb{R}^n$ is a finitely supported function from $\mathbb{R}^n$ to $\mathbb{R}$ such that if we define $\tilde{\mathbb{E}}_\mu f = \sum_{x \in \text{Supp}(\mu)} \mu(x) \cdot f(x)$ then $\tilde{\mathbb{E}}_\mu 1 = 1$ and $\tilde{\mathbb{E}}_\mu f^2 \geq 0$ for every $f \in \mathbb{R}[n]_{d/2}$. We call $\tilde{\mathbb{E}}_\mu f$ the *pseudo-expectation* of f with respect to μ . We will sometimes use the notation $\tilde{\mathbb{E}}_{\mu(x)} f(x)$ to emphasize that we apply the pseudo expectation to the polynomial f that is taken with respect to the formal variables x .

If $q \in \mathbb{R}[x_1, x_2, \dots, x_n]_{d'}$, we say that a degree d pseudo distribution μ *satisfies* the constraint $\{q \geq 0\}$ if $\tilde{\mathbb{E}}_\mu q \cdot f^2 \geq 0$ for every $f \in \mathbb{R}[n]_{(d-d')/2}$. We say that μ satisfies the constraint $\{q = 0\}$ if $\tilde{\mathbb{E}}_\mu q f = 0$ for every $f \in \mathbb{R}[x_1, x_2, \dots, x_n]_d$. We say that $\mu(x)$ is a pseudo-distribution over the sphere or the unit ball if it satisfies $\{\|x\|^2 = 1\}$ or $\{\|x\|^2 \leq 1\}$.

If μ is a degree d pseudo-distribution and $r \in \mathbb{R}[x_1, x_2, \dots, x_n]_k$ a sum-of-squares polynomial with $k \leq d$, then the degree $d-k$ pseudo distribution $\mu' = r \cdot \mu$ is called a *degree- k reweighting* of μ . Note that μ' satisfies all constraints of degree at most $d-k$ that are satisfied by μ .

The sos algorithm is given as input a set of constraints $\mathcal{E}$, a polynomial q , and a parameter d , and runs in time $n^{O(d)}$ and outputs the degree d pseudo-distribution μ that satisfies all the constraints in $\mathcal{E}$ maximizes $\tilde{\mathbb{E}}_\mu q$ (see [BKS16]). In the special case of scalar real valued random variables, there's an actual distribution that agrees with a degree d pseudo-distribution on all degree at most $d-1$ polynomials.

FACT 3.2 (SEE COROLLARY 6.14 IN [REZ00], SEE ALSO [LAS15]). *Suppose μ is a pseudo-distribution on $\mathbb{R}$ of degree d . Then, there's an actual distribution μ' over $\mathbb{R}$ such that $\mathbb{E}_{\mu'} p = \tilde{\mathbb{E}}_\mu p$ for every polynomial p of degree at most $d-1$.*

4 THE ALGORITHM

We now describe our algorithm, and show its analysis. A crucial tool for the analysis is the following general *structure theorem* on distributions over rank one matrices:

THEOREM 4.1 (STRUCTURE THEOREM FOR PSEUDO-DISTRIBUTIONS ON RANK ONE). *Let $\varepsilon > 0$, let μ be a pseudo-distribution over $\mathbb{S}^{n-1} \times \mathbb{S}^{n-1}$ of degree at least $k+2$, where $k = \sqrt{n}(\log n)^C / \varepsilon^3$ for an absolute constant $C \geq 1$. Then, μ has a degree- k reweighting μ' such that for $u_0 = \mathbb{E}_{\mu'(u,v)} u$ and $v_0 = \mathbb{E}_{\mu'(u,v)} v$,*

$$\left\| u_0 v_0^\top - \tilde{\mathbb{E}}_{\mu'(u,v)} uv^\top \right\|_F \leq \varepsilon \cdot \|u_0 v_0^\top\|_F.$$

Furthermore, we can find the reweighting polynomial $p = \mu'/\mu$ in time $2^{O(k)}$ and p has only rational coefficients in the monomial basis with numerators and denominators of magnitude at most $2^{O(k)}$.

Theorem 4.1 is proven in Section 5. Our algorithm uses it as follows:

Input: Subspace $\mathcal{W} \subseteq \mathbb{R}^{n^2}$ (in the form of a basis), and parameter $\varepsilon > 0$.

Operation:

- (1) Let $k = \tilde{O}(\sqrt{n}/\text{poly}(1/\varepsilon))$ be set as in the statement of Theorem 4.1.
- (2) Run the sum-of-squares algorithm to obtain a degree $k+2$ pseudo-distribution μ over pairs of vectors $(u, v) \in \mathbb{R}^{2n}$ that satisfies the constraint $uv^\top \in \mathcal{W}$ and $\|u\|^2 = \|v\|^2 = 1$. If no such pseudo distribution exists, output FAIL.
- (3) Use the procedure of Theorem 4.1 to find a degree k reweighting μ' of μ such that $\|\tilde{\mathbb{E}}_{\mu'} uv^\top - u_0 v_0^\top\|_F \leq \varepsilon \|u_0 v_0^\top\|_F$ where $u_0 = \tilde{\mathbb{E}}_{\mu'} u$ and $v_0 = \tilde{\mathbb{E}}_{\mu'} v$.
- (4) Output $u_0 v_0^\top$.

As discussed in Section 2.1, the following theorem immediately implies our main result (Theorem 1.3):

THEOREM 4.2 (ANALYSIS OF ALGORITHM). *Let $\varepsilon > 0$ and $\mathcal{W} \subseteq \mathbb{R}^{n^2}$ be a linear subspace. Then on input a basis for $\mathcal{W}$, if there exists a nonzero rank one matrix $uv^\top \in \mathcal{W}$ then Algorithm 4.1 will output a nonzero rank one matrix L such that $\|\Pi_{\mathcal{W}} L\|_F^2 \geq (1 - \varepsilon^2) \|L\|_F^2$ where $\Pi_{\mathcal{W}}$ is the projector to $\mathcal{W}$.*

PROOF. Under the assumptions of the theorem, there exists a nonzero rank one matrix $uv^* \in \mathcal{W}$ and by scaling we can assume $\|u\| = \|v\| = 1$ and hence the degree d SOS algorithm will return a pseudo-distribution μ satisfying these constraints for every d . Since a reweighting μ' of a pseudo-distribution μ satisfies all constraints μ satisfies, we get that $\tilde{\mathbb{E}}_{\mu'} uv^* \in \mathcal{W}$. Hence the ℓ_2 (i.e. Frobenius) distance between the nonzero rank one $u_0 v_0^*$ output by Algorithm 4.1 and the subspace $\mathcal{W}$ will be at most $\varepsilon \|u_0 v_0^*\|_F$ thus completing the proof. Note that the proof would have gone through even if the pseudo-distribution μ did not satisfy the condition that $uv^\top \in \mathcal{W}$ but merely that $\|\Pi_{\mathcal{W}^\perp} uv^\top\| \ll \|u_0 v_0^\top\|$ where Π_V is the projector to a subspace V . The proof of Theorem 4.1 actually guarantees that $\|u_0 v_0^\top\| \geq k/n$ which means that it suffices that $\|\Pi_{\mathcal{W}} uv^\top\|^2 \geq 1 - k^2/n^2$ hence implying that the proof works for the near perfect completeness case, as mentioned in Remark 1.4. $\square$

5 STRUCTURE THEOREM

In this section, we prove that every (pseudo) distribution over the unit sphere has a $\tilde{O}(\sqrt{n})$ -degree reweighting with second moment close to rank-1 in Frobenius norm. As discussed in Section 2.2, this theorem can be thought of as an approximate and robust variant of Lovett's rectangle lemma [Lov14].

THEOREM 5.1 (STRUCTURE THEOREM, REAL SYMMETRIC VERSION). *Let $\varepsilon > 0$, let μ be a pseudo-distribution over $\mathbb{S}^{n-1}$ of degree at least*

$k + 2$, where $k = \sqrt{n}(\log n)^C / \varepsilon^3$ for an absolute constant $C \geq 1$. Then, μ has a degree- k reweighting ("symmetric rank 1 reweighting") μ' such that for $m = \mathbb{E}_{\mu'} x$,

$$\left\| mm^\top - \tilde{\mathbb{E}}_{\mu'(x)} xx^\top \right\|_F \leq \varepsilon \cdot \|mm^\top\|_F.$$

Furthermore, we can find the reweighting polynomial $p = \mu' / \mu$ in time $2^{O(k)}$ and p has only rational coefficients in the monomial basis with numerators and denominators of magnitude at most $2^{O(k)}$.

Remark 5.2. Our techniques extend to show similar structure theorem for pseudo-distributions over rank $r > 1$. For e.g., in Section C of the Appendix, we give a higher-rank version of the structure theorem here.

The following more general version (see Section A for a proof) will be useful for the analysis of our algorithm from the previous section. We note that the previous theorem suffices for the symmetric analog of Algorithm 4.1.

THEOREM 5.3 (HIGHER DIMENSION STRUCTURE THEOREM). *Let $\varepsilon > 0$, let μ be a pseudo-distribution over $(u_1, u_2) \sim (\mathbb{S}^{n-1})^2$ of degree at least $k + 2$, where $k = \sqrt{n}(\log n)^C / \varepsilon^3$ for an absolute constant $C \geq 1$. Then, μ has a degree- k reweighting ("asymmetric rank 1 reweighting") μ' such that for each $1 \leq j \leq 2$*

$$\left\| m_i m_i^\top - \tilde{\mathbb{E}}_{\mu'(u_i)} uu_i^\top \right\|_F \leq \varepsilon \cdot \|m_i m_i^\top\|_F,$$

where $m_i = \mathbb{E}_{\mu'(u_i)} u_i$. Furthermore, we can find the reweighting polynomial $p = \mu' / \mu$ in time $2^{O(k)}$ and p has only rational coefficients in the monomial basis with numerators and denominators of magnitude at most $2^{O(k)}$.

Theorem 5.3 directly implies Theorem 4.1. Indeed, if we write $u = u_0 + u'$ and $v = v_0 + v'$ where u', v' are mean zero random variables, then we see that

$$\mathbb{E}(u_0 + u')(v_0 + v')^\top = u_0 v_0^\top + \mathbb{E} u' v'^\top$$

but $\|\mathbb{E} u' v'^\top\|^2 \leq \|\mathbb{E} u' u'^\top\| \|\mathbb{E} v' v'^\top\|$.

We present the proof of Theorem 5.3 which is similar to that of Theorem 5.1 in Section A of the Appendix.

5.1 Reweighting Schemes

The proof of Theorem 5.1 is based on the following general results about existence of low-degree SoS reweighting schemes. We prove these results in the following sections.

The first lemma shows that there's a low-degree reweighting for any pseudo-distribution over an interval in $\mathbb{R}$ such that the resulting distribution is concentrated around the old standard deviation.

LEMMA 5.4 (SCALAR FIXING REWEIGHTING: FIXING A SCALAR AROUND ITS STANDARD DEVIATION). *Let μ be a pseudo-distribution over $\mathbb{R}$ satisfying $\{X \leq n\}$ and $\tilde{\mathbb{E}}_\mu X^2 \geq 1$. Then, for some absolute constant C , there exists a reweighting μ' of μ of degree $k \log^C(n) / \varepsilon^2$ satisfying $\tilde{\mathbb{E}}_\mu (x - m)^k \leq \varepsilon^k m^k$ for some m satisfying $|m| \geq 1$.*

We consider this "fixing" the distribution, since if μ and μ' were actual distributions, the conclusion of Lemma 5.4 would imply that $\mathbb{P}_{\mu'}[|x - m| \geq 2\varepsilon m] \leq 2^{-k}$. Since the proof of Lemma 5.4 is rather technical, we defer it to Appendix 6.

Next, we show that distributions over the d -dimensional unit ball have $O(d)$ -degree reweightings such that the resulting distribution is concentrated around a single vector. Furthermore, the proof of this result also extends to pseudo-distribution of degree at least $O(d)$. This result is related to previous results on using high-degree sum-of-squares relaxations for optimizing general polynomials over the unit sphere [DW12]. However, the previously known bounds are not strong enough for our purposes.

LEMMA 5.5 (SUBSPACE FIXING REWEIGHTING: FIXING A DISTRIBUTION IN A SUBSPACE). *For every $C \geq 1$ there is some C' , such that if μ is a distribution over the unit ball $\{x : \|x\| \leq 1\}$ of $\mathbb{R}^d$ such that $\mathbb{E}_\mu \|x\|^2 \geq d^{-C}$ then there is a degree $k = d \cdot (\log d)^{C'}$ reweighting μ' of μ such that*

$$\left\| \mathbb{E}_{\mu'(x)} x \right\|^2 \geq 0.99 \mathbb{E}_{\mu(x)} \|x\|^2.$$

Further, the reweighting polynomial $p = \mu' / \mu$ can be found in time $2^{O(k)}$, has all coefficients upper bounded by $2^{O(k)}$ in the monomial basis, and satisfies $p(x) \leq k^{O(k)} \|x\|^k$. Moreover, the conclusions above hold as long as μ is a pseudo-distribution of degree at least $k + 2$.

We prove Lemma 5.5 in Section 7. One can use the above fixing to obtain the following corollary that allows us to grow the norm of the expectation using a low-degree reweighting.

COROLLARY 5.6. *Let μ be a distribution on $\mathbb{S}^{n-1}$ with $x_0 = \mathbb{E}_\mu x$ satisfying $\|x_0\|^2 \leq 0.5$. Then, for an absolute constant C , there's a SoS polynomial p of degree $k = \frac{\sqrt{n}}{\delta^2} \log^C(n)$ such that the reweighting $\mu' = \mu \cdot p$ satisfies: $\|\mathbb{E}_{\mu'} x\|^2 > (1 - \delta)\|x_0\|^2 + 0.5 \|\mathbb{E}_\mu x^\perp (x^\perp)^\top\|$ where $x^\perp$ is the component of x orthogonal to x_0 and is defined by $\frac{x - \langle x, x_0 \rangle x_0}{\|x_0\|^2}$ if $x_0 \neq 0$ and x otherwise. Moreover, the above conclusion holds as long as μ is a pseudo-distribution of degree at least $k + 2$.*

5.2 Proof of Structure Theorem

We now prove Theorem 5.1 using Lemmas 5.4, 5.5 and Corollary 5.6.

PROOF OF THEOREM 5.1. Our proof is based on taking iterative reweightings of the pseudo-distribution μ : we start from $\mu_0 = \mu$ with covariance $M_0 = \tilde{\mathbb{E}}_\mu [xx^\top]$ and construct $\mu_1, \mu_2, \dots, \mu_q$ in the course of our construction with means $m_t = \tilde{\mathbb{E}}_{\mu_t} x$. Let x_{-t} defined by $x - \frac{\langle x, m_t \rangle}{\|m_t\|^2} m_t$ if $m_t \neq 0$ and x otherwise, denote the component of x orthogonal to m_t and set $C_t = \tilde{\mathbb{E}}_{\mu_t} x_{-t} x_{-t}^\top$ to be the covariance matrix of x_{-t} .

We first describe the first step in the process above - where we obtain μ_1 from μ_0 . Let $\lambda_1 = \lambda_1(C_1) \geq \lambda_2 = \lambda_2(C_1) \geq \dots \geq \lambda_n = \lambda_n(C_1) \geq 0$ denote the eigenvalues of C_1 . Since μ is a distribution on $\mathbb{S}^{n-1}$, $\sum_{j=1}^n \lambda_j = \text{Tr}(C) = \tilde{\mathbb{E}}_\mu \|x\|^2 = 1$. Thus if we let $\ell = \lceil 2\sqrt{n} \rceil$,

then since the sum of the first ℓ eigenvalues is larger than any block of ℓ eigenvalues, we get that $\lceil \frac{n}{\ell} \rceil \sum_{j=1}^{\ell} \lambda_j \geq 1$ or

$$\mathbb{E}_{\mu_0} \|x_S\|^2 = \sum_{j=1}^{\ell} \lambda_j \geq 2/\sqrt{n}$$

where S is the ℓ -dimensional subspace spanned by the eigenvectors of C_0 corresponding to $\lambda_1, \dots, \lambda_{\ell}$.

Applying Lemma 5.5 to the distribution of x_S , we obtain a reweighting μ_1 of μ such that

$$\left\| \mathbb{E}_{\mu_1} x \right\|^2 \geq \left\| \mathbb{E}_{\mu_1} x_S \right\|^2 \geq 0.5 \left\| \mathbb{E}_{\mu_0} x x^T \right\| \geq \frac{1}{\sqrt{n}}. \quad (5.1)$$

Next, we describe an iterative step in the process. Suppose our current distribution and the mean are μ_t and m_t respectively for $t \geq 1$. We can assume that

$$\sum_{i=2}^n \lambda_i(C_t)^2 > \varepsilon^2 \|m_t\|^4 \quad (5.2)$$

because otherwise, we are done.

Applying Corollary 5.6 to the distribution μ_t with $\delta = \varepsilon/10$ we obtain a distribution μ_{t+1} such that:

$$\|m_{t+1}\|^2 = \left\| \mathbb{E}_{\mu_{t+1}} x \right\|^2 \geq (1 - \frac{\varepsilon}{10}) \|m_t\|^2 + 0.5\varepsilon \|m_t\|^2 \geq (1 + \frac{\varepsilon}{4}) \|m_t\|^2. \quad (5.3)$$

Now, (5.1) yields that $\|m_1\|^2 \geq \frac{1}{\sqrt{n}}$ and (5.3) yields that the $\|m_{t+1}\|^2 \geq (1 + \frac{\varepsilon}{4}) \|m_t\|^2$. Further, since at all times μ_i is a distribution on the sphere, $\|m_t\|^2 \leq 1$. Thus, the number of iterations t cannot exceed $c \log(n)/\varepsilon$ for some absolute constant c .

Finally, observe that if $p_1, p_2, \dots, p_t$ are the SoS polynomials used in the reweightings in each of the iterations above, then, the final distribution $\mu_t = \mu \cdot p$ for $p = \prod_{i \leq t} p_i$. The degree of p is thus upper bounded by $\frac{\sqrt{n}}{\varepsilon^3} \cdot \log^C(n)$ for some absolute constant C . $\square$

6 FIXING SCALAR-VALUED RANDOM VARIABLES

In this section, we prove Lemma 5.4. We begin by restating it.

LEMMA 6.1 (SCALAR FIXING REWEIGHTING). *Let μ be a pseudo-distribution over $\mathbb{R}$ satisfying $\{x \leq n\}$ and $\mathbb{E}_{\mu} x^2 \geq 1$. Then, for some absolute constant C , there exists a reweighting ("scalar fixing reweighting") μ' of μ of degree $k = Cd \log(n)/\varepsilon^2$ satisfying $\mathbb{E}_{\mu'}(x - m)^d \leq \varepsilon^d m^d$ for some m satisfying $|m| \geq 1$.*

It is instructive to derive intuition from a *conditioning* version of the lemma above for actual probability distributions. Given a random variable x with distribution μ that has standard deviation 1 and is bounded in $[-n, n]$, we know that with probability at least $\Theta(\frac{\delta}{n^2})$ probability that $x^2 \geq 1 - \delta$. As a result, the probability of at least one of $x \geq 1 - \Theta(\delta)$ or $x \leq -(1 - \Theta(\delta))$, say the former, is also at least $\Theta(\frac{\delta}{n^2})$. Next, we partition $[1 - \delta, n]$ into $O(\log(n))$ intervals with end points differing by a multiplicative factor of, say 1.1. Then,

from the above calculation, there's an interval in this partition such that x is contained in it with probability at least $\Theta(\frac{\delta}{n^2 \log(n)})$. Thus, if we condition on x lying in the above chosen interval to obtain μ' , then $KL(\mu || \mu') \leq O(\log(n) + \log(1/\delta))$.

Our plan is to roughly implement the above conditioning argument for pseudo-distributions. This demands that instead of conditioning, we use reweightings by low-degree SoS polynomials and that further, all our arguments should hold for low-degree pseudo-distributions with degree roughly matching the KL-divergence bound above.

We will use a general trick in order to aid us in this task. Specifically, Fact 3.2 says that any statement about expectations of degree d polynomials apply to all pseudo-distributions over $\mathbb{R}$ with degree at least $d + 1$. We will rely on this fact heavily in what follows. We begin with a simple claim that we will use repeatedly in what follows.

LEMMA 6.2. *Let μ be a pseudo-distribution over $\mathbb{R}$ and let x be the associated random variable. Let μ' be obtained by reweighting μ using $x^{2\ell}$ for some $\ell \in \mathbb{N}$. Then, $\mathbb{E}_{\mu'} x^2 \geq \mathbb{E}_{\mu} x^2$ so long as μ has degree at least $2\ell + 3$.*

PROOF. Since the claim is about a pseudo-distribution over $\mathbb{R}$, we can appeal to Fact 3.2 and thus it suffices to show the result for arbitrary actual probability distributions over $\mathbb{R}$. An application of Holder's inequality shows that $\mathbb{E}_{\mu'} x^2 \mathbb{E}_{\mu'} x^{2\ell} \leq \mathbb{E}_{\mu'} x^{2\ell+2}$. Rearranging yields, $\mathbb{E}_{\mu'} x^2 = \mathbb{E}_{\mu} x^{2\ell+2} / \mathbb{E}_{\mu} x^{2\ell} \geq \mathbb{E}_{\mu} x^2$. $\square$

Our main technical tool is the following lemma that shows that if we take two reweightings μ_1, μ_2 of a distribution $\mu = \mu_0$ associated with a random variable x over $\mathbb{R}$ such that the means $\mathbb{E}_{\mu_i} x^2$ remain roughly the same, then, under μ_1 , x^2 is concentrated around its mean under μ_1 . Notice that this is a statement about expectations of polynomials over *arbitrary* distributions over the real line and will thus immediately extend to pseudo-distributions using Fact 3.2 as described above.

LEMMA 6.3. *Fix $\varepsilon > 0$ and $d \in \mathbb{N}$. Let $\mu = \mu_0$ be a pseudo-distribution over $\mathbb{R}$ satisfying $\mathbb{E}_{x \sim \mu} x^2 = 1$. Then, for any $k > 4 + \frac{2d \log(1/\varepsilon)}{\varepsilon}$, and for successive reweightings $\mu_1 = \mu \cdot p_1$ and $\mu_2 = \mu_1 \cdot p_2$ of μ using the polynomials $p_1 = \frac{x^{2k}}{\mathbb{E}_{\mu_0} x^{2k}}$ and $p_2 = \frac{x^{2k}}{\mathbb{E}_{\mu_1} x^{2k}}$ respectively, at least one of the following three consequences holds:*

- (1) $\mathbb{E}_{\mu_1} x^2 > (1 + \varepsilon) \cdot \mathbb{E}_{\mu} x^2$.
- (2) $\mathbb{E}_{\mu_2} x^2 > (1 + \varepsilon) \cdot \mathbb{E}_{\mu_2} x^2$.
- (3) $\mathbb{E}_{\mu_1} (x^2 - m)^{2d} \leq 3\varepsilon^{2d} m^{2d}$ for $m = \mathbb{E}_{\mu_1} x^2$.

Moreover, the claim holds also for pseudo-distributions μ of degree at least $5k$.

Remark 6.4. Observe the three statements above are claims about expectations of degree at most $4k + 2$ polynomials under $\mu_0 = \mu$. Specifically, the three conditions have the following equivalent form:

(1)

$$\mathbb{E}_\mu x^{2d+2} > (1 + \varepsilon) \mathbb{E}_\mu x^{2d} \mathbb{E}_\mu x^2.$$

(2)

$$\mathbb{E}_\mu x^{4d+2} > (1 + \varepsilon) \mathbb{E}_\mu x^{4d} \mathbb{E}_\mu x^2.$$

(3)

$$\mathbb{E}_\mu x^{2d} \left(x^2 - \frac{\mathbb{E}_\mu x^{2d+2}}{\mathbb{E}_\mu x^{2d}} \right)^{2d} \leq 3\varepsilon^{2d} \left(\frac{\mathbb{E}_\mu x^{2d+2}}{\mathbb{E}_\mu x^{2d}} \right)^{2d} \mathbb{E}_\mu x^{2d}.$$

PROOF. Observe that by an application of Holder's inequality, $(\mathbb{E}_\mu x^{2k})^{2k} \geq (\mathbb{E}_\mu x^2)^{4k-2} \geq 1$. A similar argument shows that $\mathbb{E}_{\mu_1} x^{2k} = \frac{\mathbb{E}_\mu x^{4k}}{\mathbb{E}_\mu x^{2k}} \geq 1$. Thus, p_1, p_2 are well-defined. Now, suppose 1) and 2) do not hold. We write

$$\begin{aligned} \mathbb{E}_{\mu_1} \left(\frac{x^2}{m} - 1 \right)^{2d} &\leq 16^{2d} \varepsilon^{2d} + \mathbb{E}_{\mu_1} \left(\frac{x^2}{m} - 1 \right)^{2d} \mathbf{1} \left(\frac{x^2}{m} > (1 + \varepsilon)^4 \right) \\ &\quad + \mathbb{E}_{\mu_1} \left(\frac{x^2}{m} - 1 \right)^{2d} \mathbf{1} \left(\frac{x^2}{m} < (1 - \varepsilon)^4 \right). \end{aligned} \quad (6.1)$$

We bound the second term on the right hand side next. When 1) and 2) do not hold, observe that $\mathbb{E}_{\mu_1} x^2, \mathbb{E}_{\mu_2} x^2 \in (1 \pm \varepsilon) \mathbb{E}_\mu x^2$. Using Lemma 6.2 and the form of p_2 , this yields that for every $\ell \leq k$, $\mathbb{E}_{\mu_1} x^{2\ell+2} \leq (1 + \varepsilon) \mathbb{E}_{\mu_1} x^{2\ell} \mathbb{E}_{\mu_1} x^2$. Repeatedly applying this inequality yields:

$$\mathbb{E}_{\mu_1} x^{2k+2} \leq (1 + \varepsilon)^k \left(\mathbb{E}_{\mu_1} x^2 \right)^{k+1}. \quad (6.2)$$

Using Markov's inequality along with (6.2) yields:

$$\mathbb{P}_{\mu_1} \left[\frac{x^2}{m} \geq q \right] \leq \left(\frac{1 + \varepsilon}{q} \right)^k. \quad (6.3)$$

Let $y = \frac{x^2}{m} - 1$. Then, we have:

$$\begin{aligned} \mathbb{E}_{\mu_1} \left(\frac{x^2}{m} - 1 \right)^{2d} \mathbf{1} \left(\frac{x^2}{m} > (1 + \varepsilon)^4 \right) &\leq \int_{y \geq 1 + (1 + \varepsilon)^4} y^{2d} \mu(y) dy \\ &\leq \int_{y \geq 1 + (1 + \varepsilon)^4} y^{2d} \left(\frac{1 + \varepsilon}{1 + y} \right)^k dy \\ &\leq \int_{y \geq 1 + (1 + \varepsilon)^4} \frac{y^{2d}}{(1 + y)^{k/2}} \left(\frac{1 + \varepsilon}{\sqrt{1 + y}} \right)^k dy \\ &\leq \int_{y \geq 1 + (1 + \varepsilon)^4} \frac{y^{2d}}{(1 + y)^{k/2}} dy \\ &= \int_{y \geq 1 + (1 + \varepsilon)^4} \frac{1}{(1 + y)^2} \cdot \frac{y^{2d}}{(1 + y)^{k/2-2}} dy \end{aligned} \quad (6.4)$$

(6.5)

Now, since $k > 4 + \frac{2d}{\varepsilon}$, $y \leq \varepsilon \frac{k-4}{2d} y \leq \varepsilon(1 + y)^{\frac{k-4}{2d}}$. And thus, $\frac{y^{2d}}{(1 + y)^{k/2-2}} \leq \varepsilon^{2d}$ for every $y \geq 1 + (1 + \varepsilon)^4$. Thus, the expression in (6.4) is upper bounded by $\varepsilon^{2d} \int_{y \geq 1} \frac{1}{(1 + y)^2} \leq \varepsilon^{2d}$. This shows that the second term in (6.1) is upper bounded by ε^{2d} .

Analyzing the third term in (6.1) is easy: we have: $\mathbb{P}_{\mu_1} [x^2 < q] \leq \mathbb{P}_{\mu_0} [x^2 < q] \cdot \max_{x^2 < q} p_1(x) \leq q^{2k}$. For any $q < 1 - \varepsilon$, this quantity is at most ε^{2d} if $k > \frac{2d}{\varepsilon} \log(1/\varepsilon)$ and as a result, the third term in (6.1) is at most ε^{2d} .

This completes the proof. $\square$

Remark 6.5. It is important to note that even though *our arguments* in the proof above require higher degree polynomials ($> 5k$) - such as when we apply Holder's inequality - the statements themselves are about non-negativity of polynomials of degree at most $4k + 2$. Thus, an application of Fact 3.2 shows that these non-negativity statements, when true, hold for any pseudo-distribution of degree $\geq 4k + 3$. In particular, in situations as in the proof above, we do not have to be judicious in the use of the degree.

LEMMA 6.6. Let μ be a pseudo-distribution over $\mathbb{R}$ satisfying $\{x^2 \leq n\}$ and $\tilde{\mathbb{E}}_\mu x^2 \geq 1$. For some absolute constant C , there's a reweighting μ' of μ of degree $k = Cd \log(n)/\varepsilon^2$ such that $\tilde{\mathbb{E}}_{\mu'}(x^2 - m)^{2d} \leq \varepsilon^{2d} m^{2d}$ for $m \geq 1$.

PROOF. We take a sequence of reweightings $\mu_0 = \mu, \mu_1, \mu_2, \dots, \mu_r$ of μ such that $\mu_i/\mu_{i-1} = \frac{x^{2k}}{\mathbb{E}_{\mu_{i-1}} x^{2k}}$ for each $r \geq i \geq 1$. Observe that Lemma 6.2 implies that means of x^2 under μ_i grow monotonically and thus are all at least 1.

We then apply Lemma 6.3 to every 3-tuple $\mu_{i-1}, \mu_i, \mu_{i+1}$ for $1 \leq i \leq r - 1$. If conclusion 3) from the statement of Lemma 6.3 does not hold, then, then in every consecutive triple of reweightings $\mu_{i-1}, \mu_i, \mu_{i+1}$ as above, at least one of the consecutive pairs has a multiplicative gap of $(1 + \varepsilon)$ in the means of x^2 . Since μ satisfies $\{x^2 \leq n\}$, the number of successive triples of reweightings that do not satisfy the condition 3) of Lemma is at most $O(\log(n)/\varepsilon)$.

Thus, if we choose $r = O(\log(n)/\varepsilon)$, there's a consecutive triple of reweightings of μ , say $\mu_{i-1}, \mu_i, \mu_{i+1}$ that satisfy conclusion 3) in Lemma 6.3 for $m = \tilde{\mathbb{E}}_{\mu_i} x^2 \geq 1$. $\square$

We are now ready to prove Lemma 6.1.

PROOF OF LEMMA 6.1. $x \sim \mu$ satisfies $\{x^2 \leq n^2\}$ and $\tilde{\mathbb{E}}_\mu x^2 \geq 1$. We reweight by $x^{2\ell}$ where $\ell = O(d \log(n)/\varepsilon)$ is obtained from Lemma 6.6. Let $m \geq 1$ be such that the reweighted distribution μ' satisfies: $\tilde{\mathbb{E}}_{\mu'}(x^2 - m)^{2d} \leq \varepsilon^{2d} m^{2d}$. Or, $\tilde{\mathbb{E}}_{\mu'}(x - \sqrt{m})^{2d}(x + \sqrt{m})^{2d} \leq \varepsilon^{2d} m^{2d}$.

Now, since for every x , $(x - \sqrt{m})^{2d} + (x + \sqrt{m})^{2d} \geq \sqrt{m}^{2d}$, $\mathbb{E}(x - \sqrt{m})^{2d} + \mathbb{E}(x + \sqrt{m})^{2d} \geq \sqrt{m}^{2d}$ for every distribution over $\mathbb{R}$. Thus, for every distribution over $\mathbb{R}$, one of $\mathbb{E}(x - \sqrt{m})^{2d}$ and $\mathbb{E}(x + \sqrt{m})^{2d}$ is at least $0.5\sqrt{m}^{2d}$. Using Fact 3.2, the same conclusion holds for every pseudo-distribution of degree at least $2d + 1$ and in particular, for μ' .

Thus, say $\tilde{\mathbb{E}}_\mu(x - \sqrt{m})^{2d} \geq 0.5\sqrt{m}^{2d}$. Then, reweighting μ' by $(x - \sqrt{m})^{2d}$ to obtain μ'' yields the $\tilde{\mathbb{E}}_{\mu''}(x + \sqrt{m})^{2d} = \frac{\tilde{\mathbb{E}}_{\mu'}(x - \sqrt{m})^{2d}(x + \sqrt{m})^{2d}}{\tilde{\mathbb{E}}_{\mu'}(x - \sqrt{m})^{2d}}$ and $\tilde{\mathbb{E}}_{\mu''}(x + \sqrt{m})^{2d} \tilde{\mathbb{E}}_{\mu'}(x - \sqrt{m})^{2d} \leq$

$\varepsilon^{2d}(\sqrt{m})^{2d}$ proving the claim for $m' = -\sqrt{m}$ which is at least 1 in magnitude. $\square$

7 FIXING VECTOR-VALUED RANDOM VARIABLES

We show that distributions over the d -dimensional unit ball have $O(d)$ -degree reweightings such that the resulting distribution is concentrated around a single vector. Furthermore, the proof of this result also extends to pseudo-distribution of degree at least $O(d)$.

LEMMA 7.1 (SUBSPACE FIXING REWEIGHTING, LEMMA 5.5, RE-STATE). *Let μ be a distribution over the unit ball of $\mathbb{R}^d$. Let $C \geq 1$ be such that $\mathbb{E}_\mu \|x\|^2 \geq d^{-C}$. Then, μ has a degree- k reweighting ("subspace fixing reweighting"), where $k = d \cdot (\log d)^{C'}$ for $C' = C'(C)$ only a function of C , such that*

$$\left\| \mathbb{E}_{\mu'(x)} x \right\|^2 \geq 0.99 \mathbb{E}_{\mu(x)} \|x\|^2.$$

Further, the reweighting polynomial $p = \mu'/\mu$ can be found in time $2^{O(k)}$, has all coefficients upper bounded by $2^{O(k)}$ in the monomial basis, and satisfies $p(x) \leq k^{O(k)} \|x\|^k$. Moreover, the conclusions above hold even if μ is a pseudo-distribution of degree at least $k + 2$.

On a high level, the proof goes as follows: The final reweighting is a combination of three reweightings. The first reweighting approximately fixes the scalar variable $\|x\|^2$ as in the previous section. The second reweighting ensures that a single direction captures the expected norm in the sense that for some unit vector v the variable $\langle v, x \rangle^2$ has expectation close the expectation of $\|x\|^2$ (which also means that the second moment is close to rank-1 in trace norm). This step is the key innovation of this section. The final step is to fix the variable $\langle v, x \rangle$ such that its expectation is approximately fixed to at least the square root of the expectation of $\langle v, x \rangle^2$, which ensures that the norm of the expectation of v is large.

PROOF. Let $\varepsilon > 0$ be a sufficiently small constant and let μ be a pseudo-distribution over the unit ball of $\mathbb{R}^d$ that satisfies the requirements of the theorem. We will first find a reweighting of μ such that $\tilde{\mathbb{E}}_{\mu(x)} \langle v, x \rangle^2 \geq (1 - \varepsilon) \tilde{\mathbb{E}}_{\mu(x)} \|x\|^2$. Then, by fixing the scalar variable $\langle v, x \rangle$ as in the previous section, we obtain another reweighting that satisfies the requirements of the theorem.

By fixing the scalar variable $\|x\|^2$ as in the previous section, we may assume that μ satisfies $\tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k} \leq (1 + \varepsilon)^k \cdot (\tilde{\mathbb{E}}_{\mu(x)} \|x\|^2)^k$ for all $k \leq (\varepsilon^{-1} \log d)^{O(1)}$. By Hölder's inequality for pseudo-expectations (this follows by an application of Fact 3.2 to the scalar random variable $\|x\|^2$), we also have $\tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k+2} \geq \tilde{\mathbb{E}}_{\mu(x)} \|x\|^2 \cdot \tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k}$.

Let v be a random unit vector in $\mathbb{R}^d$. It is known that for every $k \in \mathbb{N}$ and $x \in \mathbb{R}^n$, $\mathbb{E}_v \langle v, x \rangle^{2k} = c_k \|x\|^{2k}$ where c_k is the k -th raw moment of the Beta distribution with parameter 1 and $d - 2$ [Sta82]. If $k \leq O(d/\varepsilon)$ then these moments satisfy that $c_{k+1} \geq (1 - \varepsilon)c_k$ and $1 \geq c_k \geq k^{-O(k)}$.

Using the bound $c_{k+1} \geq (1 - \varepsilon)c_k$ and the fact that the variable $\|x\|^2$ is approximately fixed, it follows that

$$\mathbb{E}_{\frac{\tilde{\mathbb{E}}}{v \mu(x)} \langle v, x \rangle^{2k}} = c_k \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{2k}, \quad (7.1)$$

$$\mathbb{E}_{\frac{\tilde{\mathbb{E}}}{v \mu(x)} \langle v, x \rangle^{2k+2}} = c_{k+1} \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{2k+2} \quad (7.2)$$

$$\geq (1 - \varepsilon)^2 c_k \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^2 \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{2k}. \quad (7.3)$$

Combining the above two equations,

$$\mathbb{E}_{\frac{\tilde{\mathbb{E}}}{v \mu(x)} \langle v, x \rangle^{2k+2}} \geq (1 - \varepsilon)^2 \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^2 \cdot \mathbb{E}_{\frac{\tilde{\mathbb{E}}}{v \mu(x)} \langle v, x \rangle^{2k}}. \quad (7.4)$$

which means that there exists a unit vector v such that $\tilde{\mathbb{E}}_{\mu(x)} \langle v, x \rangle^{2k+2} \geq (1 - \varepsilon)^2 \tilde{\mathbb{E}}_{\mu(x)} \|x\|^2 \cdot \mathbb{E}_v \tilde{\mathbb{E}}_{\mu(x)} \langle v, x \rangle^{2k}$. Therefore, the reweighting polynomial $x \mapsto \langle v, x \rangle^{2k} / \tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k}$ yields a reweighting μ' such that $\tilde{\mathbb{E}}_{\mu'(x)} \langle v, x \rangle^2 \geq (1 - \varepsilon)^2 \cdot \tilde{\mathbb{E}}_{\mu(x)} \|x\|^2$.

It remains to argue that we can find such a unit vector in time $2^{O(k)}$. To that end we will show that a random unit vector succeeds with probability $2^{-O(k)}$. That argument will also show that the normalization factor $\tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k}$ is bounded by $k^{O(k)}$.

The key step is to bound the variance of the variable $\tilde{\mathbb{E}}_{\mu(x)} \langle v, x \rangle^{2k}$ (over the randomness of v).

$$\mathbb{E}_v \left(\frac{\tilde{\mathbb{E}}}{\mu(x)} \langle v, x \rangle^{2k} \right)^2 \leq \mathbb{E}_v \frac{\tilde{\mathbb{E}}}{\mu(x)} \langle v, x \rangle^{4k} = c_{2k} \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{4k} \quad (7.5)$$

$$\leq k^{O(k)} \cdot \left(c_k \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{2k} \right) \quad (7.6)$$

The first inequality is Cauchy-Schwarz for pseudo-expectations. The second inequality uses the bounds $c_{2k} \leq k^{O(k)} c_k^2$ and $\tilde{\mathbb{E}}_{\mu(x)} \|x\|^{4k} \leq (1 + \varepsilon)^{2k} (\tilde{\mathbb{E}}_{\mu(x)} \|x\|^{2k})^2$ (the latter holds because we approximately fixed the scalar variable $\|x\|^2$).

A standard Markov-like inequality (see for example [BKS15, Lemma 5.3]) shows that the following event over random unit vectors v has probability at least $k^{-O(k)}$ (note that this probability is w.r.t. the distribution of the random variable v , which is an actual distribution),

$$\left\{ \begin{array}{l} \frac{\tilde{\mathbb{E}}}{\mu(x)} \langle v, x \rangle^{2k+2} \geq (1 - \varepsilon)^3 \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^2 \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \frac{\tilde{\mathbb{E}}}{\mu(x)} \langle v, x \rangle^{2k} \\ \frac{\tilde{\mathbb{E}}}{\mu(x)} \langle v, x \rangle^{2k} \geq k^{O(-k)} \cdot \frac{\tilde{\mathbb{E}}}{\mu(x)} \|x\|^{2k} \end{array} \right\}$$

Any unit vector that satisfies the above conditions yields a reweighting polynomial $p(x) \propto \langle v, x \rangle^{2k}$ that satisfies the conclusion of the theorem. $\square$

We now move on to showing Corollary 5.6 restated here.

COROLLARY 7.2. *Let μ be a distribution on $\mathbb{S}^{n-1}$ with $m_0 = \mathbb{E}_\mu x$ satisfying $\|m_0\|^2 \leq 0.5$. Then, for an absolute constant C , there's a SoS polynomial p of degree $k = \frac{\sqrt{n}}{\delta^2} \log^C(n)$ such that the reweighting $\mu' = \mu \cdot p$ satisfies: $\|\mathbb{E}_{\mu'} x\|^2 > (1 - \delta) \|m_0\|^2 + 0.5 \|\mathbb{E}_\mu x_\perp x_\perp^\top\|$ where $x_\perp$ is the component of x orthogonal to m_0 and is defined by $\frac{x - \langle x, m_0 \rangle m_0}{\|m_0\|^2}$.*

if $m_0 \neq 0$ and x otherwise. Moreover, the above conclusion holds also for all pseudo-distributions of degree at least $k + 2$.

PROOF. The key idea of the proof is to observe that there is a subspace S , orthogonal to m_0 , of roughly $\sqrt{n}$ dimensions such that $\mathbb{E}_\mu \|x_S\|^2 \geq \|\mathbb{E}_\mu x_\perp\|^2$. In fact, one can take S to be the explicit subspace spanned by the top $\approx \sqrt{n}$ eigenvectors of $\mathbb{E}_\mu x_\perp x_\perp^\top$.

Once we find such an S , we can apply the subspace fixing reweighting (Lemma 7.1) to x_S , the projection of x to the subspace S . After an appropriate change of basis, this is a random variable in $\mathbb{R}^{\sqrt{n}}$. Thus, Lemma 7.1 gives us a $\approx \sqrt{n}$ degree reweighting that ensures that $\|\mathbb{E}_\mu x_S\|^2 \geq 0.5 \|\mathbb{E}_\mu x_S x_S^\top\|$ and by virtue of being in S , this is orthogonal to $\mathbb{E}_\mu x$, we obtain the needed gain in the norm of the expectation of x that we needed. However, to guarantee the claim in the Corollary, we also need to ensure that $\mathbb{E}_{\mu'} \langle x, \mathbb{E}_\mu x \rangle^2 \geq (1 - \delta) \mathbb{E}_{\mu'} \langle x, \mathbb{E}_\mu x \rangle^2$. reweighting μ could a priori lead to a situation where this is no longer true. To “protect” the norm of the previous expectation, we “fix” the scalar random variable $\langle x, \mathbb{E}_\mu x \rangle$ using scalar fixing reweighting (Lemma 6.1) before applying the subspace fixing reweighting (Lemma 7.1). Our final reweighting is a combination of the above two reweightings.

We now formalize the above idea. First, $\mathbb{E}_\mu \|x\|^2 \geq 1 - \|m_0\|^2 \geq 0.5$. We first reweight the scalar random variable $\langle x, \mathbb{E}_\mu x \rangle^2$ using the scalar fixing reweighting (Lemma 6.1) for some d to be chosen later and ensure:

$$\mathbb{E}_{\mu_1(x)} \left(\langle x, \mathbb{E}_\mu x \rangle - \|\mathbb{E}_\mu x\|^2 \right)^{4d} \leq \varepsilon^{4d} \|\mathbb{E}_\mu x\|^{8d}. \quad (7.7)$$

Let μ_1 be the resulting distribution. Observe that using Lemma 6.2, $\mathbb{E}_{\mu_1} \langle x, \mathbb{E}_\mu x \rangle^2$ is not smaller than $\mathbb{E}_\mu \langle x, \mathbb{E}_\mu x \rangle^2 = \|\mathbb{E}_\mu x\|^2$. If $\mathbb{E}_{\mu_1} \langle x, \mathbb{E}_\mu x \rangle^2 \geq 0.9$, we are done. Thus, assume that $\|\mathbb{E}_{\mu_1} x\| \geq 0.1$.

Let $C = \mathbb{E}_{\mu_1} x_\perp x_\perp^\top$ with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$. For $\ell = \lceil \sqrt{n} \rceil + 1$, let S be the subspace of the eigenvectors of C corresponding to eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_\ell$. For any $x \in \mathbb{R}^n$, let x_S denote the projection of x to the subspace S . Our plan is to apply the subspace fixing reweighting (Lemma 7.1) to the random variable x_S that lives in $\mathbb{R}^\ell$. For this, we need to lower bound $\mathbb{E}_{\mu_1} \|x_S\|^2$.

We proceed to do that next. Observe that

$$\sum_{i=\ell+1}^n \lambda_i^2 \leq n \lambda_{\ell+1}^2 \leq \frac{n}{\ell(\ell-1)} \sum_{i \neq j} \lambda_i \lambda_j. \quad (7.8)$$

Thus,

$$\begin{aligned} \left(\mathbb{E}_{\mu_1} \|x_S\|^2 \right)^2 &= \left(\sum_{i=1}^{\ell} \lambda_i \right)^2 \\ &\geq \sum_i \lambda_i^2 + \sum_{i \neq j \in [\ell]} \lambda_i \lambda_j \\ &\geq \sum_{i \leq \ell} \lambda_i^2 + \frac{\ell(\ell-1)}{n} \sum_{i=\ell+1}^n \lambda_i^2 \end{aligned}$$

$$\begin{aligned} &\geq \sum_{i=1}^n \lambda_i^2 \\ &= \left\| \mathbb{E}_{\mu_1} x_\perp x_\perp^\top \right\|^2. \end{aligned} \quad (7.9)$$

Next, we reweight x_S using the subspace fixing reweighting (Lemma 7.1) to obtain μ_2 .

Observe that $\mathbb{E}_{\mu_2} x_S \in S$ and $\mathbb{E}_\mu x$ is orthogonal to S . Thus, $\|\mathbb{E}_{\mu_2} x\|^2 \geq \frac{\mathbb{E}_{\mu_2} \langle x, \mathbb{E}_{\mu_1} x \rangle^2}{\|\mathbb{E}_\mu x\|^2} + \|\mathbb{E}_{\mu_2} x_S\|^2$. Thus, to complete the proof,

we must analyze two quantities: $\frac{\mathbb{E}_{\mu_2} \langle x, \mathbb{E}_{\mu_1} x \rangle^2}{\|\mathbb{E}_\mu x\|^2}$ and $\|\mathbb{E}_{\mu_2} x_S\|^2$. The latter is guaranteed to be at least $0.5 \cdot \|\mathbb{E}_{\mu_1} x_\perp x_\perp^\top\|$ using (7.9) and Lemma 7.1. We now analyze $\mathbb{E}_{\mu_2} \langle x, \mathbb{E}_{\mu_1} x \rangle^2$.

Let $\mu_2/\mu_1 = R$ for the polynomial R defining the subspace fixing reweighting (obtained in Lemma 7.1) satisfying $\mathbb{E}_{\mu_1(x)} R(x) = 1$ and

$$|R(x)| \leq n^{O(\sqrt{n} \log(n))} \|x\|^{O(\sqrt{n} \log(n))}, \quad (7.10)$$

for every x .

We thus have:

$$\begin{aligned} \left(\mathbb{E}_{\mu_2(x)} \left(\langle x, \mathbb{E}_\mu x \rangle - \|\mathbb{E}_\mu x\|^2 \right)^2 \right)^{4d} &\leq \left(\mathbb{E}_{\mu_2(x)} \left(\langle x, \mathbb{E}_\mu x \rangle - \|\mathbb{E}_\mu x\|^2 \right)^2 \right)^{2d} \\ &= \frac{\left(\mathbb{E}_{\mu_1(x)} R(x) \left(\langle x, \mathbb{E}_\mu x \rangle - \|\mathbb{E}_\mu x\|^2 \right)^2 \right)^{2d}}{\mathbb{E}_{\mu_1(x)} R(x)} \\ &\leq \mathbb{E}_{\mu_1(x)} R(x)^2 \cdot \mathbb{E}_{\mu_1(x)} \left(\langle x, \mathbb{E}_\mu x \rangle - \|\mathbb{E}_\mu x\|^2 \right)^{4d} \\ &\leq n^{O(\sqrt{n} \log(n))} \cdot \mathbb{E}_{\mu_1(x)} \|x\|^{O(\sqrt{n} \log(n))} \varepsilon^{4d} \|\mathbb{E}_\mu x\|^{8d}. \end{aligned}$$

The first inequality uses Holder's inequality for pseudo-distributions, the second inequality uses Cauchy-Schwarz inequality for pseudo-distributions and the final inequality uses the bound (7.7). The equality in step 2 uses the definition that $R = \mu_2/\mu_1$.

For $\varepsilon \leq \text{poly}(\delta)$ and $d = O(\sqrt{n} \log(n) \log(1/\delta))$, we can ensure that the final expression above is at most δ^{8d} and completing the proof. $\square$

8 CONCLUSIONS AND FURTHER DIRECTIONS

We have shown an $\exp(\tilde{O}(\sqrt{n}))$ time algorithm for the $\text{BSS}_{1,1-\varepsilon}$ problem, or equivalently to the problem of finding (up to small ℓ_2 error) a rank one matrix in a subspace. This raises several questions. One question, mentioned in Remark 1.4 is whether we can remove the perfect, or near-perfect, completeness condition. Another, possibly more important, question is whether one can improve the exponent $\sqrt{n}$ further. Indeed, it has been suggested that the $\text{BSS}_{c,s}$ problem for constants $s < c$ might be in quasi-polynomial time (which is closely related to the conjecture that $\text{QMA}(2) \subseteq \text{EXP}$). Given that our algorithm is inspired by Lovett's $\tilde{O}(\sqrt{n})$ bound on the communication complexity of rank n matrices, it is tempting to speculate that the full log rank conjecture (i.e., a $\text{polylog}(n)$ bound)

would imply such a quasi-polynomial time algorithm. We think a black box reduction from the log rank conjecture to such an algorithm is unlikely. For starters, we would need the proof of the log rank conjecture to be embedded in the sos proof system. But even beyond that, it seems that we need more general statements, that (unlike the log rank conjecture) do not apply only to *Boolean* matrices. There do seem to be natural such statements that could imply improved algorithmic results. In particular, we believe resolving the following two questions could help in making such progress:

Question 8.1. Is it the case that for every distribution μ over $\mathbb{S}^{n-1}$ and every $\varepsilon, \delta > 0$ there is a (not necessarily positive) function $r : \mathbb{S}^{n-1} \rightarrow \mathbb{R}$ such that $\mathbb{E}_{v \sim \mu} |r(v)| = 1$, $\mathbb{E}_{v \sim \mu} |r(v)| \log |r(v)| \leq O(n^\delta)$ and a nonzero rank one L such that

$$\| \mathbb{E}_{v \sim \mu} [r(v) v v^\top] - L \|_F \leq \varepsilon \|L\|_F ?$$

A positive solution for Question 8.1 for any $\delta < 1/2$ would be very interesting. It may⁸ improve the best known bound for the log rank conjecture to $\tilde{O}(n^\delta)$ and if appropriately extended to pseudo-distributions, improve our algorithm's running time to $\exp(\tilde{O}(n^\delta))$ as well. We do know that the answer to this question is *No* if one does not allow *negative* reweighting functions.

Another interesting question is the following:

Question 8.2. Is there a function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for every $\delta > 0$ and a distribution μ over $\mathbb{S}^{n-1}$, there is an $O(n^\delta)$ round reweighting μ' of μ such that

$$\mathbb{E}_{v, v' \sim \mu} \langle v, v' \rangle^4 = f(\delta) \left(\mathbb{E}_{v, v' \sim \mu} \langle v, v' \rangle^2 \right)^2 ?$$

We do not know of a way to use a positive answer for Question 8.2 for an improved bound on the log rank conjecture, but (an appropriate sos version of) it does imply an improved algorithms for the problem of “2 vs 4 provers QMA” where in the completeness case there are four non-entangled provers, while the soundness condition holds even if, say, provers 1 and 2, as well as provers 3 and 4, share an entangled state. This variant seems “morally related” to the *small set expansion* problem (though we do not know of a reduction in either direction).

A PROOF OF THEOREM 5.3

PROOF OF THEOREM 5.3. Let μ be the pseudo-distribution on $(u_1, u_2) \in (\mathbb{S}^{n-1})^2$. As in the proof of Theorem 5.1, our final reweighting is obtained by combining a sequence of reweightings $\mu_0 = \mu, \mu_1, \mu_2, \dots, \mu_T$. Let m_1^j, m_2^j denote $\tilde{\mathbb{E}}_{\mu_j} u_1, \tilde{\mathbb{E}}_{\mu_j} u_2$, respectively.

In the first step, for each $1 \leq i \leq 2$, we do the following:

- (1) Let S_i be the subspace spanned by the largest $\lceil 2\sqrt{n} \rceil$ eigenvectors of $\mathbb{E}_{\mu(u_i)} u_i u_i^\top$. reweight $\mu(u_i)$ in the subspace S_i using Lemma 5.5 using $\tilde{O}(\sqrt{n})$ degree.

⁸As mentioned in Footnote 7, improving the bounds on the log rank conjecture might require better control of the dependence of the bound on ε than we need for our setting.

- (2) reweight $\langle u_i, m_i \rangle$ using Lemma 5.4 using $\tilde{O}(\sqrt{n}/(\varepsilon)^2)$ -degree.

Similar to the proof of Corollary 5.6, we can argue that at the end of the first step, $\|m_1\|^2 \|m_2\|^2 \geq \Theta(1/n)$.

In any iterative step, we do the following. If there's an i such that $\|m_i m_i^\top - \mathbb{E}_{\mu(u_i)} u_i u_i^\top\|_F > \varepsilon \|u_i u_i^\top\|$,

- (1) Let S_i be the subspace spanned by the largest $\lceil 2\sqrt{n} \rceil$ eigenvectors of $\mathbb{E}_{\mu(u_i)} u_i^\perp u_i^{\perp\top}$ where $u_i^\perp$ is the component of u_i orthogonal to m_i . reweight $\mu(u_i)$ in the subspace S_i using Corollary 5.6 using $\tilde{O}(\sqrt{n}/\varepsilon^2)$ degree.
- (2) reweight $\langle u_i, m_i \rangle$ using Lemma 5.4 using $\tilde{O}(\sqrt{n}/\varepsilon^2)$ -degree.

We now track the potential function $\|m_1\|^2 \|m_2\|^2$. In any step, the second reweighting above implies that under any reweighting $\|m_i\|$ doesn't decrease by a factor of more than $1 - \varepsilon/10$. The first reweighting yields that at least one of $\|m_1\|^2$ or $\|m_2\|^2$ increases by a factor of $(1 + \varepsilon/2)$. In effect, after each reweighting, the potential rises by a multiplicative $(1 + \Theta(\varepsilon))$. Since $\|m_1\|^2 \|m_2\|^2 \leq 1$ and at least $\Theta(1/n)$ after the first step, the number of steps in the reweighting is upper bounded by $O(\log(n)/\varepsilon)$ giving the result. $\square$

B REDUCTION BETWEEN REAL AND COMPLEX BEST SEPARABLE STATE PROBLEMS

LEMMA B.1. For every subspace $\mathcal{W} \subseteq \mathbb{C}^{n^2}$, there's a subspace $\mathcal{Y} \subseteq \mathbb{R}^{4n^2}$ such that:

- (1) *Completeness:* If there's an $x, y \in \mathbb{S}^{n-1}(\mathbb{C})$ such that $xy^* \in \mathcal{W}$, then there's a $u, v \in \mathbb{S}^{2n-1}$ such that $uv^* \in \mathcal{Y}$.
- (2) *Soundness:* If there's a $U \in \mathcal{Y}$ and u_0, v_0 such that $\|u_0 v_0^\top - U\|_F \leq \varepsilon \|u_0 v_0^\top\|_F$, then there's a $X \in \mathcal{W}$ and an x_0, y_0 such that $\|x_0 y_0^* - X\|_F \leq \varepsilon \|x_0 y_0^*\|_F$.

PROOF. It is easiest to describe the construction of the subspace $\mathcal{U}$ from $\mathcal{W}$ in two steps. Let $\langle W^j, X \rangle = 0$ for $j \leq \text{codim}(\mathcal{W})$ be the linear constraints that define $\mathcal{W}$. Write $X = A + iB$ for $i = \sqrt{-1}$ and $W^j = C^j + iD^j$. Then, $X \in \mathcal{W}$ iff for every $j \leq \text{codim}(\mathcal{W})$,

$$\begin{cases} \langle C^j, A \rangle + \langle D^j, B \rangle = 0 \\ \langle D^j, A \rangle - \langle C^j, B \rangle = 0. \end{cases}$$

$\square$

Let $\mathcal{W}' \subseteq \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n}$ be the subspace of ordered pairs of $n \times n$ matrices (A, B) satisfying (B) for each $j \leq \text{dim}(\mathcal{W})$.

Observe that by construction, a matrix $X = A + iB \in \mathcal{W}$ iff the pair $(A, B) \in \mathcal{W}'$. Next, we define a subspace $\mathcal{Y} \subseteq \mathbb{R}^{2n \times 2n}$ as follows. We think of each $Y \in \mathcal{Y}$ as a 2×2 block matrix of $n \times n$ matrices with the blocks being labeled as $Y_{11}, Y_{12}, Y_{21}, Y_{22}$ in the natural way. We define $\mathcal{Y}$

$$Y \in \mathcal{Y} \text{ iff } (Y_{11} + Y_{22}, Y_{21} - Y_{12}) \in \mathcal{W}'. \quad (\text{B.1})$$

We now claim that the subspace $\mathcal{Y}$ satisfies the requirements of the Lemma. First observe that if $Y \in \mathcal{Y}$, then by our construction,

$(Y_{11} + Y_{22}, Y_{21} - Y_{12}) \in \mathcal{W}'$ and consequently,

$$X = (Y_{11} + Y_{22}) + i(Y_{21} - Y_{12}) \in \mathcal{W}'. \quad (\text{B.2})$$

Completeness. If $xy^* \in \mathcal{W}$ then, writing $x = u + iv$ and $y = u' + iv'$ and setting $A = uu^\top + vv^\top$ and $B = vu^\top - uv^\top$ yields that $(A, B) \in \mathcal{W}'$ and thus, consequently, $Y = (u, v)(u', v')^\top \in \mathcal{Y}$.

Soundness. Suppose $Y \in \mathcal{Y}$ and there's u, v such that $\|uv^\top - Y\|_F \leq \varepsilon \|uv^\top\|$. Let u_1, u_2 (v_1, v_2) be the components of u in the first and second column (row) blocks respectively. From (B.2), we know that $X = A + iB$ for $A = (Y_{11} + Y_{22}) + i(Y_{21} - Y_{12}) \in \mathcal{W}'$. Let $U = u_1 + iu_2$ and $V = v_1 + iv_2$. Then, we can rewrite the above as:

$$X = A + iB = (u_1 + iu_2)(v_1 + iv_2)^* + (Y_{11} + Y_{22} - u_1v_1^\top - u_2v_2^\top) + i(Y_{21} - Y_{22} - u_2v_2^\top + u_1v_1^\top) \in \mathcal{W}'. \quad (\text{B.3})$$

Now, $\|(u_1 + iu_2)(v_1 + iv_2)^*\|^2 = \|u_1\|^2 + \|u_2\|^2 + \|v_1\|^2 + \|v_2\|^2$. And by an application of triangle inequality,

$$\begin{aligned} & \|(Y_{11} + Y_{22} - u_1v_1^\top - u_2v_2^\top) + i(Y_{21} - Y_{22} - u_2v_2^\top + u_1v_1^\top)\| \\ & \leq \sum_{s,t=1}^2 \|Y_{st} - u_s v_t^\top\| \\ & \leq \|uv^\top - Y\|_F \leq \varepsilon \|uv^\top\| \\ & \leq \varepsilon \|U\| \|V\| \\ & = \varepsilon \|UV^*\|. \end{aligned}$$

C HIGHER RANK STRUCTURE THEOREM

THEOREM C.1. *Let $\varepsilon > 0$, let μ be a pseudo-distribution over $(u_1, u_2, \dots, u_r)$ such that $\sum_i \|u_i\|^2 = 1$. Let the degree of μ be at least $k + 2$, where $k = \sqrt{rn}(\log n)^C / \varepsilon^3$ for an absolute constant $C \geq 1$. Then, μ has a degree- k reweighting μ' such that for each $1 \leq j \leq r$*

$$\sum_{i=1}^r \left\| m_i m_i^\top - \mathbb{E}_{\mu'(u_i)} uu^\top \right\|_F^2 \leq \varepsilon^2 \cdot \left(\sum_{i=1}^r \|m_i m_i^\top\|_F^2 \right),$$

where $m_i = \mathbb{E}_{\mu'(u_i)} u_i$. Furthermore, we can find the reweighting polynomial $p = \mu' / \mu$ in time $2^{O(k)}$ and p has only rational coefficients in the monomial basis with numerators and denominators of magnitude at most $2^{O(k)}$.

Remark C.2. The above theorem can be used to design an algorithm that finds symmetric rank r matrices of Frobenius norm 1 inside subspaces $\mathcal{W}$ in time $2^{\tilde{O}(\sqrt{rn}/\varepsilon^2)}$.

PROOF. For every $(u_1, u_2, \dots, u_r)$ in the support of μ , define $u \in \mathbb{R}^{rn}$ be the concatenation of $u_1, u_2, \dots, u_r$. Then, μ can be equivalently thought of as a distribution over the unit sphere. Then, $\|u\|^2 = 1$. By Theorem 5.1, there's a $\tilde{O}(\sqrt{rn}/\varepsilon^2)$ degree reweighting μ' of μ such that there exists a u^0 such that $\|u^0 u^{0\top} - \mathbb{E}_{\mu'} uu^\top\|_F \leq \varepsilon \|u^0 u^{0\top}\|_F$.

Let $m_i = \mathbb{E}_{\mu'}(u^0)_i (u^0)_i^\top$. Then, $\sum_{i=1}^r \|m_i\|^2 = \|u^0\|^2$. On the other hand, $\|\mathbb{E}_{\mu'} uu^\top - \sum_{i=1}^r m_i m_i^\top\|_F^2 \leq \|uu^\top - u^0 u^{0\top}\|_F^2 \leq \varepsilon^2 \|u^0 u^{0\top}\|_F^2 = \varepsilon^2 \sum_{i=1}^r \|m_i\|^2$.

This completes the proof. $\square$

ACKNOWLEDGEMENT

We thank the anonymous reviewers for suggestions on improved presentation of the paper.

REFERENCES

- [BaCY11] Fernando G.S.L. Brandão, Matthias Christandl, and Jon Yard, *A quasipolynomial-time algorithm for the quantum separability problem*, Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '11, ACM, 2011, pp. 343–352.
- [BBH⁺12] Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou, *Hypercontractivity, sum-of-squares proofs, and their applications*, STOC, ACM, 2012, pp. 307–326.
- [BCHW16] Fernando GSL Brandao, Matthias Christandl, Aram W Harrow, and Michael Walter, *The mathematics of entanglement*, arXiv preprint arXiv:1604.01790 (2016).
- [BH15] Fernando G. S. L. Brandão and Aram Wettroth Harrow, *Estimating operator norms using covering nets*, CoRR **abs/1509.05065** (2015).
- [BKS15] Boaz Barak, Jonathan A. Kelner, and David Steurer, *Dictionary learning and tensor decomposition via the sum-of-squares method*, STOC, ACM, 2015, pp. 143–151.
- [BKS16] Boaz Barak, Praveesh Kothari, and David Steurer, *Proofs, beliefs, and algorithms through the lens of sum-of-squares*, 2016, Lecture notes in preparation, available on <http://sumofsquares.org>.
- [BT09] Hugue Blier and Alain Tapp, *All languages in np have very short quantum proofs*, Quantum, Nano and Micro Technologies, 2009. ICQNM'09. Third International Conference on, IEEE, 2009, pp. 34–37.
- [DPS04] Andrew C Doherty, Pablo A Parrilo, and Federico M Spedalieri, *Complete family of separability criteria*, Physical Review A **69** (2004), no. 2, 022308.
- [DW12] Andrew C. Doherty and Stephanie Wehner, *Convergence of SDP hierarchies for polynomial optimization on the hypersphere*, CoRR **abs/1210.5048** (2012).
- [Gha10] Sevag Gharibian, *Strong np-hardness of the quantum separability problem*, Quantum Information & Computation **10** (2010), no. 3, 343–360.
- [GL14] Dmitry Gavinsky and Shachar Lovett, *En route to the log-rank conjecture: New reductions and equivalent formulations*, ICALP (1), Lecture Notes in Computer Science, vol. 8572, Springer, 2014, pp. 514–524.
- [Gur03] Leonid Gurvits, *Classical deterministic complexity of edmonds' problem and quantum entanglement*, Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, ACM, 2003, pp. 10–19.
- [HHH96] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki, *Separability of mixed states: necessary and sufficient conditions*, Physics Letters A **223** (1996), no. 1, 1–8.
- [HM13] Aram Wettroth Harrow and Ashley Montanaro, *Testing product states, quantum merlin-arthur games and tensor optimization*, J. ACM **60** (2013), no. 1, 3.
- [Las01] Jean B. Lasserre, *An explicit exact SDP relaxation for nonlinear 0-1 programs*, IPCO, Lecture Notes in Computer Science, vol. 2081, Springer, 2001, pp. 293–303.
- [Las15] Jean Bernard Lasserre, *An introduction to polynomial and semi-algebraic optimization*, no. 52, Cambridge University Press, 2015.
- [LKCH00] Maciej Lewenstein, B Kraus, JI Cirac, and P Horodecki, *Optimization of entanglement witnesses*, Physical Review A **62** (2000), no. 5, 052310.
- [Lov14] Shachar Lovett, *Communication is bounded by root of rank*, STOC, ACM, 2014, pp. 842–846.
- [LS88] László Lovász and Michael E. Saks, *Lattices, möbius functions and communication complexity*, FOCS, IEEE Computer Society, 1988, pp. 81–90.
- [NW94] Noam Nisan and Avi Wigderson, *On rank vs. communication complexity*, FOCS, IEEE Computer Society, 1994, pp. 831–836.
- [Par00] Pablo A Parrilo, *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*, Ph.D. thesis, Citeseer, 2000.
- [Rez00] Bruce Reznick, *Some concrete aspects of hilbert's 17th problem*, Contemporary Mathematics **253** (2000), 251–272.
- [Rot14] Thomas Rothvoß, *A direct proof for lovett's bound on the communication complexity of low rank matrices*, CoRR **abs/1409.6366** (2014).

- [Sta82] A. J. Stam, *Limit theorems for uniform distributions on spheres in high-dimensional Euclidean spaces*, J. Appl. Probab. **19** (1982), no. 1, 221–228. MR 644435
- [Ved08] Vlatko Vedral, *Quantifying entanglement in macroscopic systems*, Nature **453** (2008), no. 7198, 1004–1007.
- [Vid03] Guifré Vidal, *Efficient classical simulation of slightly entangled quantum computations*, Phys. Rev. Lett. **91** (2003), 147902.