



Matrix Discrepancy from Quantum Communication

Samuel B. Hopkins
 Massachusetts Institute of
 Technology
 USA

Prasad Raghavendra
 University of California, Berkeley
 USA

Abhishek Shetty
 University of California, Berkeley
 USA

ABSTRACT

We develop a novel connection between discrepancy minimization and (quantum) communication complexity. As an application, we resolve a substantial special case of the *Matrix Spencer* conjecture. In particular, we show that for every collection of symmetric $n \times n$ matrices A_1, \dots, A_n with $\|A_i\| \leq 1$ and $\|A_i\|_F \leq n^{1/4}$ there exist signs $x \in \{\pm 1\}^n$ such that the maximum eigenvalue of $\sum_{i \leq n} x_i A_i$ is at most $O(\sqrt{n})$. We give a polynomial-time algorithm based on partial coloring and semidefinite programming to find such x .

Our techniques open a new avenue to use tools from communication complexity and information theory to study discrepancy. The proof of our main result combines a simple compression scheme for transcripts of repeated (quantum) communication protocols with quantum state purification, the Holevo bound from quantum information, and tools from sketching and dimensionality reduction. Our approach also offers a promising avenue to resolve the Matrix Spencer conjecture completely – we show it is implied by a natural conjecture in quantum communication complexity.

CCS CONCEPTS

• Theory of computation → Generating random combinatorial structures; Discrete optimization; Communication complexity; • Mathematics of computing → Combinatorics.

KEYWORDS

Matrix Discrepancy, Quantum Communication Complexity, Quantum Random Access Codes, Semidefinite Programming, Sketching

ACM Reference Format:

Samuel B. Hopkins, Prasad Raghavendra, and Abhishek Shetty. 2022. Matrix Discrepancy from Quantum Communication. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3519935.3519954>

1 INTRODUCTION

In this paper we study *discrepancy minimization* for matrices. To set up our main problem, let us begin with the classic result of Spencer, “six standard deviations suffice.” Let $v_1, \dots, v_n \in \mathbb{R}^m$ have $\|v_i\|_\infty \leq 1$. The goal is to assign signs $x_1, \dots, x_n \in \{\pm 1\}$ to the vectors so as to minimize $\|\sum_{i \leq n} x_i v_i\|_\infty$. As a shorthand, we often call the latter quantity the discrepancy of x . For some intuition, note that if



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

STOC '22, June 20–24, 2022, Rome, Italy

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9264-8/22/06.

<https://doi.org/10.1145/3519935.3519954>

the vectors $v_i \in \{0, 1\}^m$ then, treating them as incidence vectors, they define a set system with n atoms and m subsets. The goal then becomes to assign x_1, \dots, x_n so as to minimize the maximum difference between the number of $+1$ ’s and -1 ’s in each set.

Choosing x_1, \dots, x_n at random presents a natural benchmark – in this case, $\mathbb{E}_{x \sim \{\pm 1\}^n} \|\sum x_i v_i\|_\infty \leq O(\sqrt{n \log m})$, by a Chernoff/union bound argument. While many similar applications of the union bound in combinatorics give tight results, Spencer’s result remarkably shows that for any v_1, \dots, v_n , this bound can in fact be beaten.

THEOREM 1.1 ([32]). *For all $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_i\|_\infty \leq 1$, there exist $x_1, \dots, x_n \in \{-1, 1\}$ such that $\|\sum_{i \leq n} x_i v_i\|_\infty \leq O(\sqrt{n \log(m/n)})$.*

In particular, if $m = O(n)$, Spencer’s result shows that a signing of discrepancy $O(\sqrt{n})$ always exists.¹ Spencer’s original result was nonconstructive, but a following a breakthrough by Bansal [4], several polynomial-time algorithms are now known to find such a signing, e.g. [14, 22, 31].

Matrix Discrepancy. We generalize the preceding setting by replacing the vectors v_1, \dots, v_n with symmetric matrices $A_1, \dots, A_n \in \mathbb{R}^{m \times m}$ having spectral norms $\|A_i\| \leq 1$.² Now the goal is to find x_1, \dots, x_n to minimize the spectral norm $\|\sum_{i \leq n} x_i A_i\|$. Note that we can recover the vector case by taking the A_i ’s to be diagonal, or more generally, commuting.

The *matrix Chernoff* bound of Ahlswede and Winter shows that, as in the vector setting, randomly choosing x gives a signing of discrepancy $\|\sum_{i \leq n} x_i A_i\| \leq O(\sqrt{n \log m})$ [1]. This inequality and its generalizations have become crucial tools in mathematics and theoretical computer science, including in applications of the probabilistic method, for instance in spectral graph theory and unsupervised learning, e.g. [18, 33]. It is a natural question to ask whether it, too, can be improved by careful choice of x – this is the content of the *Matrix Spencer* conjecture:

CONJECTURE 1.2 (MATRIX SPENCER [25, 35]). *For all $A_1, \dots, A_n \in \mathbb{R}^{m \times m}$ with $\|A_i\| \leq 1$ there exists $x \in \{\pm 1\}^n$ such that $\|\sum_{i \leq n} x_i A_i\| \leq O(\sqrt{n \log(m/n)})$.*

Despite significant effort, this conjecture has remained largely open for a decade, with partial progress in the block-diagonal and rank-one cases [21, 23, 30]. Thus, resolving Matrix Spencer (even in a substantial special case) seems likely to lead to new techniques in discrepancy.

¹And, in fact, the constant in the big- O is at most 6, hence the name.

²We expect that the main results in this paper continue to hold if \mathbb{R} is replaced by \mathbb{C} . Furthermore, if the matrices A_i are not symmetric/Hermitian, they can be replaced by their “Hermitian dilations” $\begin{pmatrix} 0 & A_i \\ A_i^\top & 0 \end{pmatrix}$ without changing any of the asymptotic bounds in this paper.

1.1 Results

We resolve the Matrix Spencer conjecture in the case that A_1, \dots, A_n have *moderate rank*. More formally, in addition to the assumption $\|A_i\| \leq 1$, we additionally assume that $\|A_i\|_F \leq n^{1/4}$, where $\|\cdot\|_F$ is the Frobenius norm.

THEOREM 1.3 (MODERATE-RANK MATRIX SPENCER). *Let $A_1, \dots, A_n \in \mathbb{R}^{m \times m}$ have $\|A_i\| \leq 1$ and $\|A_i\|_F \leq n^{1/4}$. Then there exists $x \in \{\pm 1\}^n$ such that $\|\sum_{i \leq n} x_i A_i\| \leq O(\sqrt{n \log(m/n)})$. Furthermore, such an x can be found in polynomial time.*

Even in the presence of the “moderate rank” assumption $\|A_i\|_F \leq n^{1/4}$, our result captures settings where the looser bound $O(\sqrt{n \log m})$ is un-improvable for randomly-chosen x – for instance, if the A_i ’s are all diagonal with nonzero entries in the first \sqrt{n} diagonal entries.³ Thus, our result captures a novel improvement over the matrix Chernoff bound.

To prove Theorem 1.3, we introduce a new approach to discrepancy minimization using (one-way) communication complexity. In the matrix case, this connection leads us to quantum communication. For starters, we give a new proof of Spencer’s theorem: after translation into a (classical) communication problem, Spencer’s theorem can be proved using a simple compression scheme for repeated communication protocols. To prove our moderate-rank Matrix Spencer theorem, we combine a quantum analogue of this compression scheme with several other tools, including quantum state purification, sketching/dimensionality reduction, and consequences of the *Holevo bound* from quantum information theory.

Discrepancy bounds proved using our techniques are automatically algorithmic. In the vector (Spencer) case, our arguments give a new analysis of the randomized linear programming approach first analyzed by Eldan and Singh [14]. In the matrix case, we give an analogous algorithm based on semidefinite programming. (This algorithm uses a very different semidefinite program than Bansal’s original use of semidefinite programming in the vector case.)

The $O(\sqrt{n})$ discrepancy bound is tight even for the case of rank one matrices. This is witnessed by considering the matrices $(1, e_i)(1, e_i)^\top$ where e_i are the standard basis vectors. Note that it cannot be tight for diagonal matrices under this assumption, since n vectors in \mathbb{R}^n with ℓ_2 norms $n^{1/4}$ have discrepancy at most $\tilde{O}(n^{1/4})$ (the *Komlós* setting) [3]. This suggests a number of interesting questions beyond Matrix Spencer: are there matrix analogues of other discrepancy bounds for vectors, for instance under ℓ_2 assumptions (as in the Komlós setting) or ℓ_2 and ℓ_∞ assumptions (like the Beck-Fiala setting)?

We hope that opening the way to use communication complexity techniques to prove results in discrepancy leads to future progress. As an illustration, we show that our techniques offer a promising avenue to fully resolve the Matrix Spencer conjecture – we now describe a natural conjecture in quantum communication complexity which would imply it.

To describe the conjecture we need a small amount of notation. Let $\text{index} : \{-1, 1\}^n \times [n] \rightarrow \{-1, 1\}$ be the *index function*, given by $\text{index}(x, i) = x_i$. The index function induces the following one-way communication problem between two players, Alice and Bob. Alice receives $x \in \{-1, 1\}^n$ and Bob receives $i \in [n]$. Alice sends Bob

³We thank Tselil Schramm and Boaz Barak for pointing this out.

a message a , after which Bob must output a bit $b(a, i) \in \{-1, 1\}$; their goal is to jointly compute $\text{index}(x, i)$.

The main question in one-way communication complexity is: how long must Alice’s message be? This could depend on several things:

- The nature of Alice’s message – classical or quantum.
- The probability of success $\mathbb{P}(b(a, i) = x_i)$ (where the probability is over randomness in the protocol).
- The distribution of Alice and Bob’s inputs – they could be uniformly random, worst-case, or something else.

Later, we will thoroughly discuss the one-way communication complexity of the index function, after which the following conjecture will be less mysterious. For now, we state the conjecture as an illustration of the surprising connection between discrepancy and communication.

CONJECTURE 1.4 (QUANTUM ONE-WAY COMMUNICATION COMPLEXITY IN THE SMALL-ADVANTAGE REGIME). *Suppose Alice’s message ρ consists of q qubits, and Bob has ε advantage over random guessing in computing x_i for a large set of indices i , in the following sense. For each x there is a set of coordinates $S_x \subseteq [n]$ with $|S_x| \geq (1 - \delta)n$ such that $\mathbb{E}_{x \sim \{-1, 1\}^n} \min_{i \in S_x} \mathbb{P}(b(\rho, i) = x_i) \geq 1/2 + \varepsilon$. Then for every small-enough $\delta > 0$, if $\varepsilon \geq 1/\sqrt{n}$, Alice must send $q \geq \log(1/\varepsilon^2) + \Omega(\varepsilon^2 n)$ qubits.*

Note that Conjecture 1.4 remains interesting even if $\delta = 0$; indeed, this special case is most interesting from a quantum communication point of view, and we expect that it already contains most of the challenge in proving the conjecture.

Using the same argument as for Theorem 1.3 but substituting the communication lower bound in Conjecture 1.4 for a weaker version we prove in the course of proving Theorem 1.3, our techniques show:

THEOREM 1.5. *Suppose Conjecture 1.4 is true. Then the Matrix Spencer conjecture holds, and there is a polynomial-time algorithm based on semidefinite programming to find the signing it promises.*

The classical analogue of Conjecture 1.4 is true; we record a proof in this paper, although we believe it is probably known implicitly in the literature. In fact, using our techniques, the classical analogue gives a new algorithmic proof of Spencer’s theorem. Our proof of Theorem 1.3 establishes a special case of Conjecture 1.4 where Alice must send a *pure state*, from which (with some work) we are able to deduce our moderate-rank Matrix Spencer theorem.

1.2 Techniques

1.2.1 From Discrepancy to Communication.

Discrepancy Is Exactly Average-Bob One-Way Communication Complexity. To build intuition, we start with the following simple observation. Let $R_{\text{worst,unif}} = R_{\text{worst,unif}}(n, \varepsilon)$ be the minimum length of a message a that Alice must send to Bob in a one-way classical protocol for the n -bit index function in order to achieve

$$\min_{x \in \{\pm 1\}^n} \mathbb{E}_{i \sim [n]} \mathbb{P}(b(a, i) = x_i) \geq \frac{1}{2} + \varepsilon.$$

Here, the subscript “worst,unif” denotes that Alice’s input is worst-case over $x \in \{\pm 1\}^n$ and Bob’s is uniform in $[n]$. Similarly, define

$Q_{\text{worst,unif}}$ for quantum one-way communication. The following claim shows that *lower bounds* on $R_{\text{worst,unif}}$ and $Q_{\text{worst,unif}}$ imply *upper bounds* on discrepancy for vectors and matrices.

Claim 1.6. For every integer $m > 0$, if $R_{\text{worst,unif}}(n, \varepsilon) > \log m + O(1)$ then for $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_i\|_\infty \leq 1$ there exists $x \in \{-1, 1\}^n$ such that $\|\sum_{i \leq n} x_i v_i\|_\infty \leq 2\varepsilon n$. Conversely, if $R_{\text{worst,unif}}(n, \varepsilon) < \log m - O(1)$, then there exist $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_i\|_\infty \leq 1$ such that for all $x \in \{\pm 1\}^n$ we have $\|\sum_{i \leq n} x_i v_i\|_\infty \geq 2\varepsilon n$. Furthermore, the same holds if we replace R by Q and the v s by matrices $A_1, \dots, A_m \in \mathbb{R}^{m \times m}$ with $\|A_i\| \leq 1$.

PROOF. We show one direction of the proof in the classical case; the other direction and the quantum case are similar. Suppose $v_1, \dots, v_n \in \mathbb{R}^m$ have $\|v_i\|_\infty \leq 1$ but for every $x \in \{\pm 1\}^n$ we have $\|\sum_{i \leq n} x_i v_i\|_\infty > 2\varepsilon n$. Then for each x we may associate a standard basis vector $y_x \in \mathbb{R}^m$ such that $|\langle y_x, \sum_{i \leq n} x_i v_i \rangle| > 2\varepsilon n$. This induces a $\log m + O(1)$ bit communication protocol as follows. On input x , Alice sends Bob the name of the coordinate j represented by y_x , as well as the sign s of $\langle y_x, \sum_{i \leq n} x_i v_i \rangle$. Bob outputs a biased random bit $b(j, i)$ with expectation $s \cdot v_i(j)$. Then for each x we can compute:

$$\mathbb{E}_i [\mathbb{P}(b(j, i) = x_i)] = \mathbb{E}_i \left[\frac{1}{2} + s \cdot x_i \cdot \frac{v_i(j)}{2} \right] > \frac{1}{2} + \varepsilon.$$

For the other direction, observe that a protocol for the index function where there are m possible messages Alice may send induces a set of vectors $v_1, \dots, v_n \in \{\pm 1\}^m$ by writing out Bob's outputs. The success probability of this protocol gives a lower bound on the discrepancy of v_1, \dots, v_n . \square

In spite of its simplicity, we do not know how to use this connection between discrepancy and $R_{\text{worst,unif}}$ and $Q_{\text{worst,unif}}$ to prove any interesting discrepancy upper bounds. The difficulty is that in the regime of interest for (Matrix) Spencer-type theorems, $m \approx n$ and $\varepsilon \approx 1/\sqrt{n}$. That is, Bob has very tiny advantage over oblivious random guessing in determining $\text{index}(x, i)$, and Alice is sending just a logarithmic number of (qu)bits. We are not aware of any direct techniques to lower bound $R_{\text{worst,unif}}$, let alone $Q_{\text{worst,unif}}$ in this regime. (Of course, an indirect argument is available for $R_{\text{worst,unif}}$ by appealing to Spencer's discrepancy result.)

Trading Average Bob for Average Alice. Our first key technical contribution is another connection between discrepancy and communication, but for $R_{\text{unif,worst}}$ and $Q_{\text{unif,worst}}$ rather than $R_{\text{worst,unif}}, Q_{\text{worst,unif}}$ – that is, now Alice's input will be random, but Bob's will be worst-case. While this difference may seem small, the requirement that (for typical x) Bob has nontrivial advantage over random guessing in computing x_i for all i makes it much easier to prove lower bounds – we will see why momentarily. (Actually, our lower bounds will apply even when Bob has nontrivial advantage for, say, $0.9n$ coordinates $i \in [n]$ – this technical improvement is important for the connection to discrepancy, but we will mainly ignore it for simplicity in this introduction.)

We now discuss the key lemma we prove connecting communication and *partial coloring*, starting with the following standard definition:

Definition 1.7 (Partial coloring). A partial coloring of matrices A_1, \dots, A_n with discrepancy $\Delta > 0$ is a vector $x \in [-1, 1]^n$ such

that $|x_i| = 1$ for a constant fraction of coordinates $i \in [n]$, and $\|\sum_{i \leq n} x_i A_i\| \leq \Delta$.⁴ (A similar definition applies for the case of vectors v_1, \dots, v_n .)

It is a standard result that Spencer-style discrepancy theorems can be proved by alternately finding partial colorings and removing vectors/matrices which have been fully colored (i.e., they have $|x_i| = 1$), so it suffices to prove the existence of partial colorings with small discrepancy.

For simplicity in this introduction, we restrict attention to the setting where the number of vectors/matrices is the same as the dimension – i.e. $A_1, \dots, A_n \in \mathbb{R}^{n \times n}$ or $v_1, \dots, v_n \in \mathbb{R}^n$ – in which case we are looking for partial colorings of discrepancy $O(\sqrt{n})$. And, for now, we drop algorithmic considerations and worry only about the existence of partial colorings.

LEMMA 1.8 (SPECIAL CASE OF THE COMPRESS OR COLOR LEMMA (LEMMA 3.2), INFORMAL). Suppose $A_1, \dots, A_n \in \mathbb{R}^{n \times n}$ with $\|A_i\| \leq 1$ lack a partial coloring with discrepancy $O(\sqrt{n})$. Then there is a quantum one-way communication protocol for the n -bit index function of the following form. Alice sends a $\log n + 1$ qubit message ρ . If Bob gets input i , he measures ρ in the eigenbasis of the matrix $\begin{pmatrix} A_i & 0 \\ 0 & -A_i \end{pmatrix}$, receiving an eigenvalue λ_i as an outcome; then he outputs a random bit $b(\rho, i)$ with bias λ_i . This protocol has the following guarantee: for every $x \in \{\pm 1\}^n$ there is a set $S_x \subseteq [n]$ with $|S_x| \geq 0.99n$ such that

$$\mathbb{E}_{x \in \{\pm 1\}^n} \min_{i \in S_x} \mathbb{P}(b(\rho, i) = x_i) = \frac{1}{2} + \varepsilon, \quad \varepsilon \gg \frac{1}{\sqrt{n}}.$$

From this lemma we can see the origin of Conjecture 1.4 and Theorem 1.5. It also shows that to prove our moderate-rank Matrix Spencer theorem, it suffices to rule out $\log n$ -qubit protocols for the index function with advantage $\varepsilon \gg 1/\sqrt{n}$ where Bob's measurement matrices have $\|A_i\|_F \leq n^{1/4}$.

We make a few more remarks about the Compress or Color Lemma (Lemma 3.2) before we move on to communication lower bounds, since we think the general version of the lemma is of independent interest.

General norms In full generality, the lemma says that for *any* collection of vectors v_1, \dots, v_n and *any* norm $\|\cdot\|$, either v_1, \dots, v_n admit a small- $\|\cdot\|$ -discrepancy partial fractional coloring (i.e. a coloring where $\|\sum_{i \leq n} x_i v_i\|$ is small) or v_1, \dots, v_n induce a certain kind of compression of the hypercube $\{\pm 1\}^n$ into the dual ball of $\|\cdot\|$. In the case that $\|\cdot\|$ is ℓ_∞ , this compression turns out to be a classical communication protocol. When $\|\cdot\|$ is the spectral norm, the result is a quantum communication protocol.

Rademacher width Second, the proof of the lemma goes via studying the *Rademacher width* of the set of partial colorings. The Rademacher width of a set $\mathcal{K} \subseteq \mathbb{R}^n$ is $\mathbb{E}_{g \sim \{\pm 1\}^n} \max_{x \in \mathcal{K}} \langle g, x \rangle$ – it is a standard measure of the size of \mathcal{K} . A long-established technique in discrepancy is to study the *Gaussian volume* of the partial colorings. By studying width instead, we can prove the lemma using tools from convex programming, in particular strong duality. Gaussian width was previously studied in the context of discrepancy by Eldan and Singh and by Reis and Rothvoss [14, 29]; we borrow some tools from Eldan and Singh in the proof of the Compress or Color

⁴This is often called a “fractional” partial coloring in the literature; since all partial colorings in this paper are fractional we drop the modifier.

Lemma. (The switch from Gaussian to Rademacher width – that is, using ± 1 -valued coordinates in g – is just a technical convenience.) Polynomial-time consequences Finally, since the heart of the Compress or Color Lemma is a convex program, it also has algorithmic consequences when that convex program is efficiently solvable. In particular, our proof of the contrapositive of the above statement, that communication lower bounds imply the existence of partial colorings, actually proves something stronger: such a partial coloring can be found (with high probability) by drawing a random $g \sim \{\pm 1\}^n$ and maximizing $\langle g, x \rangle$ over partial colorings x with low discrepancy. Note that this is a convex program – in particular, for the matrix discrepancy setting, it is a semidefinite program.

1.2.2 Communication Lower Bounds in the Small-Advantage Regime. Now that we have seen that communication lower bounds imply the existence of partial colorings, we need to prove some communication lower bounds.

Classical. To build some intuition, we start with the classical case. According to Lemma 1.8 (instantiated with diagonal matrices), to prove that a partial coloring of any $v_1, \dots, v_n \in \mathbb{R}^n$ with $\|v_i\|_\infty \leq O(\sqrt{n})$ exists it will suffice to rule out $\log n$ -bit one-way protocols for the n -bit index function where Alice’s input is random, Bob’s is worst case, and they have advantage $\epsilon \gg 1/\sqrt{n}$ over oblivious random guessing. (To avoid technicalities, for now we consider protocols where Bob has this advantage on all inputs $i \in [n]$, rather than just $0.99n$ of them.)

To see the subtlety of the lower bound we need to establish, let us first consider what we could get from naive information-theoretic arguments. By directly analyzing the mutual information between Alice’s input and Bob’s output, we could show that Alice must send at least $(1 - H(1/2 + \epsilon))n$ bits, where H is the binary entropy function. For small ϵ , we have $(1 - H(1/2 + \epsilon))n \approx \epsilon^2 n$ – this lower bound degrades to just $O(1)$ when $\epsilon \approx 1/\sqrt{n}$, while we need a bound larger than $\log n$.

Indeed, if Bob’s input is also random, there is actually an $O(1)$ -bit protocol achieving advantage $\epsilon \gg 1/\sqrt{n}$. Even with worst-case inputs, there is a $\log n + O(1)$ -bit protocol based on Hadamard matrices which achieves advantage $\epsilon \geq \Omega(1/\sqrt{n})$. (For both protocols, see Section ??.) This shows that we must use worst-case-ness of Bob’s input in our lower bound, and even when we do, our argument must be tight up to additive constants. We now sketch a simple argument satisfying both of these requirements.

LEMMA 1.9. *Any classical protocol for the n -bit index function achieving advantage $\epsilon \gg 1/\sqrt{n}$ when Alice’s input is uniformly random and Bob’s is worst-case requires Alice to send more than $\log n$ bits. That is, $R_{\text{unif},\text{worst}}(n, \epsilon) > \log n$ when $\epsilon \gg 1/\sqrt{n}$.*

PROOF SKETCH. Suppose for contradiction that a $\log n$ bit protocol exists with advantage $\epsilon \gg 1/\sqrt{n}$. By repeating the protocol $O(1/\epsilon^2) \ll n$ times, the players can amplify their success probability to 0.9. Concretely, in this *amplified protocol*, Alice receives x and makes $O(1/\epsilon^2)$ independent draws from the distribution over messages she would send on input x in the original protocol. She sends all of these messages to Bob, who computes all of the outputs he would compute in the original protocol and takes a majority vote. (Note that this amplification relies on Bob having a worst-case input – otherwise, Bob might already have success probability

0.9 on a few inputs and exactly $1/2$ on the rest, in which case the amplification does not have the desired effect.)

Naively, Alice is now sending $\Omega(\log n/\epsilon^2) \gg n$ bits, but we claim that her message in the amplified protocol can be compressed down to $n/10$ bits (for appropriate $\epsilon \gg 1/\sqrt{n}$). This leads to a contradiction, since she is sending Bob at least $n/2$ bits of information.

To see this, observe that it actually suffices for Alice to send a histogram of her $O(1/\epsilon^2)$ messages from the original-protocol distribution, since Bob does not need to know the ordering of the $O(1/\epsilon^2)$ messages. Since Alice’s individual messages are $\log n$ bits, there are only n possible messages in the original protocol, so she is sending multi-subset/histogram of $[n]$ of size $O(1/\epsilon^2)$. A simple counting argument shows that there are approximately $\binom{n}{O(1/\epsilon^2)}$ such histograms (this is exactly the number if there are no repeated messages, but repeated messages do not change the asymptotics). Since $1/\epsilon^2 \ll n$, Alice can now send just $\log \binom{n}{O(1/\epsilon^2)} \ll n$ bits. \square

Carrying out this argument carefully actually shows the following quantitative bounds, which may be independently interesting:

- If $\epsilon \gg 1/\sqrt{n}$, then $R_{\text{unif},\text{worst}}(n, \epsilon) \geq \log(1/\epsilon^2) + \Omega(\epsilon^2 n)$.
- If $\epsilon \ll 1/\sqrt{n}$, then $R_{\text{unif},\text{worst}}(n, \epsilon) \geq \log n - \log \log(1/\epsilon^2) - O(1)$.

This means that even for very small ϵ , like $2^{-n^{0.99}}$, Alice still must send $\Omega(\log n)$ bits. Using our discrepancy-to-communication technique, this “microscopic ϵ ” bound implies discrepancy bounds for set systems with many more atoms than sets – with n atoms and m sets, discrepancy $O(\sqrt{m})$ is achievable. (There are also generic reductions from $n \gg m$ to $n = m$, using linear programming.)

Quantum. We turn to the case of quantum communication lower bounds, which we need for the matrix discrepancy setting. To prove the (unrestricted) matrix Spencer conjecture, we would like to prove a quantum analogue of Lemma 1.9. Unfortunately, the histogram-based compression used in the argument above seems inherently classical, so another idea is needed.

To prove moderate-rank matrix Spencer, our second key technical contribution is a quantum analogue of Lemma 1.9 when Bob’s measurement matrices have $\|A_i\|_F \leq n^{1/4}$. By Lemma 1.8, this shows that partial colorings exist for every family of $A_1, \dots, A_n \in \mathbb{R}^{n \times n}$ with $\|A_i\| \leq 1$ and $\|A_i\|_F \leq n^{1/4}$.

For simplicity in this introduction, consider the case that each matrix A_i has all eigenvalues in $\{-1, 0, 1\}$, with at most \sqrt{n} nonzero eigenvalues. Alice gets a randomly chosen $x \in \{\pm 1\}^n$ and sends Bob a $\log n$ -qubit mixed state, represented by a density matrix $\rho = \rho_x$. Given any input i , Bob measures ρ_x with A_i , getting back an eigenvalue. If he receives -1 or 1 , he outputs the result; otherwise he outputs -1 or 1 uniformly at random. We want to show:

LEMMA 1.10. *Alice and Bob cannot achieve success probability $1/2 + \epsilon$ for $\epsilon \gg 1/\sqrt{n}$ by the above protocol.*

We now sketch the proof of Lemma 1.10, which takes several ingredients.

Ruling out pure-state protocols The first step is to prove a lower bound against *pure state* protocols with (potentially) *full-rank measurements*. That is, we consider the case that Alice actually sends a pure state $a_x \in \mathbb{R}^n$, and Bob is allowed to use any measurements

$M_i \in \mathbb{R}^{n \times n}$ with ± 1 eigenvalues. In this case, we can use a similar amplify-then-compress approach as in the classical case. We sketch the proof here – for details, see Theorem 5.4 for the communication lower bound and Lemma 3.4 for a version adapted to the discrepancy bound we need to prove.

In a little bit more detail, to amplify from success probability $1/2 + \varepsilon$ to 0.9, Alice sends $O(1/\varepsilon^2)$ copies of her message; Bob makes his measurement independently on each of them and takes a majority vote. Naively, this requires Alice to send around $\log n/\varepsilon^2 \gg n$ qubits, but the state that Alice sends, $a_x^{\otimes O(1/\varepsilon^2)}$, actually lies in the *symmetric subspace* of $(\mathbb{R}^n)^{\otimes O(1/\varepsilon^2)}$. This subspace, $\text{span}\{a^{\otimes O(1/\varepsilon^2)} : a \in \mathbb{R}^n\}$, has dimension roughly $(O(1/\varepsilon^2))^n$ by the same counting argument as we used in the classical case, which means that Alice can compress her message into $\log(O(n/\varepsilon^2)) \ll n$ qubits. As in the classical case, this argument crucially uses that Bob succeeds on any (worst-case) input $i \in [n]$.

Since Alice is sending a quantum state, this situation is no longer ruled out by classical information theory. However, a known consequence of the Holevo bound from quantum information says that Alice cannot communicate the n classical bits x without sending $\Omega(n)$ qubits. This is not a trivial consequence of the Holevo bound, since Bob cannot necessarily read more than one of the bits of x without collapsing the state he is sent in a way which prevents reading any of the remaining bits. This situation has been considered before, however: a result of Ambainis, Nayak, Ta-Shma, and Vazirani on *quantum random access codes* shows that it is still impossible for Alice to send $\ll n$ qubits even if Bob can only read one coordinate of x (so long as he may choose this coordinate at will) [2].

Reduction from moderate-rank to pure-state protocols We now sketch an argument that if there is a protocol of the type described in Lemma 1.10, where Alice may be sending a mixed state ρ_x , then there is also a pure-state protocol of the sort we just ruled out. We call this the *purify-then-sketch* transformation (Lemma 3.3).

First, we may assume that Alice's mixed states $\rho_x \in \mathbb{R}^{n \times n}$ have rank at most \sqrt{n} – this is because we can take them to be extremal solutions to a semidefinite program involving n linear constraints, one for each matrix A_i [9, 27]. Using this bound on the rank of the ρ_x 's, we can use quantum state purification to replace them with pure states \tilde{a}_x of dimension $\mathbb{R}^{n^{3/2}}$; when Bob measures the purified states he uses measurement matrices $A_i \otimes I$, where I is identity in \sqrt{n} dimensions. This gives a pure-state protocol with the same success probability as the protocol we started with (since the outcomes of Bob's measurements will have exactly the same distributions as before), but now Alice has to send $\frac{3}{2} \log n$ qubits, so we cannot apply the above lower bound against pure-state protocols.

To fix this, Alice replaces \tilde{a}_x with a random sketch $a_x = S\tilde{a}_x$ of it down to n dimensions, and hence $\log n$ qubits. (Here S is a random sketching matrix). Bob replaces $A_i \otimes I$ with $S(A_i \otimes I)S^\top$. We argue that *so long as A_i has at most \sqrt{n} nonzero eigenvalues* this sketching matrix preserves the success probability $1/2 + \varepsilon$ when $\varepsilon \gg 1/\sqrt{n}$. To give a tight analysis of quantities like the variance of the outcomes of the protocol after sketching – i.e., second-moment quantities like $\langle (S\tilde{a}_x)(S\tilde{a}_x)^\top, S(A_i \otimes I)S^\top \rangle^2$ – we use a combinatorial moment-method argument, which crucially uses our bounds on $\|A_i\|_F$. We also employ a number of tools from random matrix

theory – decoupling inequalities, net-based arguments, and the Hansen-Wright inequality. For details, see Section 6.

At the end, we arrive at a $\log n$ -qubit pure-state protocol with advantage $\gg 1/\sqrt{n}$, which we have already showed is impossible. This completes the proof sketch of Lemma 1.10, which in turn completes our proof sketch of the moderate-rank Matrix Spencer theorem.

1.3 Related Work

Discrepancy. Discrepancy theory is rich and well explored area of combinatorics with connections to many areas of mathematics and theoretical computer science. It has found applications in diverse areas such as approximation algorithms, differential privacy and probability theory. For a more thorough introduction, see [11, 24].

Classical results in combinatorial discrepancy are often based on linear dependencies (see [8]) and counting arguments (e.g. Beck's partial coloring method and Spencer's entropy method). Spencer's six standard deviations theorem was initially proved by combining partial coloring and counting arguments [32]. Another proof of this theorem was given by [16, 17] by making connections between existence of partial colorings and convex geometry. This theorem gave natural conditions to find partial colorings in general convex bodies.

Successful as they were, these techniques were all non-constructive and thus did not provide algorithmic insights on constructing colorings. In a breakthrough result, [5] gave the first algorithm to find the signs promised by Spencer's theorem. The algorithm was based on semidefinite programming but needed to assume the existence of a good coloring in the analysis. [22] gave a random walk-based algorithm whose analysis does not appeal to Spencer's theorem. [31] gave an elegant algorithm that produces the partial colorings in convex sets guaranteed by Gluskin's theorem. [14] provide an alternative algorithm for this problem using linear programming by providing connections to the width of the convex set – our SDP-based algorithm is a direct descendent of theirs.

Another line of work in algorithmic discrepancy is constructing algorithms for the *Beck-Fiala* and *Komlós* settings, where additional assumptions on the vectors v_1, \dots, v_n lead to tighter discrepancy bounds. Here obtaining tight bounds remains an open problem, even non-algorithmically. The best known non-constructive bounds are obtained using a technique introduced by [3] which also draws from connections to convex geometry. A recent line of work resolved the question of algorithmically matching Banaszczyk's bound [6, 7, 12]. For an overview of this line of work, see [15].

Matrix discrepancy and spectral graph theory. Another line of work that is closely related to discrepancy is the construction of sparsifiers for graphs. [10] construct linear sized (weighted) sparsifiers for graphs that approximate the Laplacian of the graph. In a celebrated work, [23] use a novel technique based on interlacing polynomials to resolve the Kadison–Singer conjecture, which can be interpreted as a tight discrepancy bound for signed sums of rank-one matrices $a_1 a_1^\top, \dots, a_n a_n^\top$ in isotropic position. This is also related to constructing unweighted sparsifiers for graphs. It remains an excellent open problem to find algorithms matching the bounds of [23].

The Matrix Spencer conjecture is a natural matrix generalization of the Spencer theorem, asking if we can improve upon the matrix Chernoff bound to get a bound similar to the one guaranteed by Spencer's theorem. It was popularized in a [blog post](#) by Raghu Meka. This bound was also conjectured in [35]. [30] provide an algorithm gives a discrepancy bound of $O(\sqrt{n \log q})$ for matrices that are block diagonal with block size q . A result by [21] uses techniques from [23] to resolve the conjecture for rank one matrices. [29] bring together techniques from the convex geometric approach of [16, 17] for constructing algorithms for graph sparsification problems akin to [10] towards potentially using these ideas to resolve the matrix Spencer conjecture.

Independent and concurrent work: In independent and concurrent work, [13] provide an different framework based on covering numbers and mirror descent to make progress on the Matrix Spencer conjecture. In the low rank setting, they show that the discrepancy $\|\sum_{i \leq n} x_i A_i\| \leq \sqrt{n \log(\min(rm/n, r))}$ where r is the rank, m is the dimension and n is the number of matrices. This bound on the rank is stronger than the one required in our result. For example, when $m = \Theta(n)$, this reduces to the constant rank case. In addition, they also show that their framework can be used to address related questions in matrix discrepancy such as the Matrix Spencer conjecture for block diagonal matrices and Schatten norms.

Communication complexity and quantum random access codes. Lower bounds in one-way communication complexity are widely used to prove lower bounds in other settings: data structures and streaming algorithms, to name just two. See [28] for a modern introduction to communication complexity. The index function, in particular, plays a central role in one-way communication, see e.g. [20]. It is a folklore result that the one-way constant-error classical communication complexity of the n -bit index function is $\Omega(n)$.

Quantum protocols for the index function also go by the name *quantum random access codes*, which have been studied intensively in the physics literature, including experimental demonstrations of quantum protocols whose success probabilities are strictly better than those achievable by classical protocols for small n , e.g. [34]. [2] show that the one-way constant-error quantum communication complexity of the n -bit index function is $\Omega(n)$. This argument was simplified and refined in [26].

2 PRELIMINARIES

For vectors $x, y \in \mathbb{R}^d$, let $\langle x, y \rangle$ denote the standard inner product $\sum_i x_i y_i$. For matrices A, B , this inner product also corresponds to $\langle A, B \rangle = \text{Tr}(AB)$. Let $\|x\|_2 = \sqrt{\sum_i x_i^2}$ denote the ℓ_2 norm, $\|x\|_\infty = \max_i |x_i|$ denote the ℓ_∞ norm and $\|x\|_1 = \sum_i |x_i|$ denote the ℓ_1 norm.

For a matrix $A \in \mathbb{R}^{d \times d}$, denote by $\|A\|_F = \sqrt{\text{Tr}(A^\top A)}$ the Frobenius norm, by $\|A\| = \sup_{\|x\|=1} \|Ax\|$ the operator or spectral norm and by $\|A\|_1 = \text{Tr}(\sqrt{A^\top A})$, the nuclear or the trace norm. For matrices A, B , $A \leq B$ if $B - A$ is a positive semidefinite matrix.

A convex set \mathcal{K} is said to be centrally symmetric if $x \in \mathcal{K}$ implies $-x \in \mathcal{K}$. For any convex set \mathcal{K} , denote by $p_{\mathcal{K}}$ the Minkowski functional defined by $p_{\mathcal{K}}(x) = \inf\{r > 0 : x \in r\mathcal{K}\}$. We say that a convex set has non-empty interior if there is an $\varepsilon > 0$ such that

for all x such that $\|x\|_2 \leq \varepsilon$, $x \in K$. A compact, convex set with non-empty interior is referred to as a convex body. If \mathcal{K} is a symmetric convex body, $p_{\mathcal{K}}$ corresponds to a norm which we denote by $\|\cdot\|_{\mathcal{K}}$. Furthermore, any norm can be seen as the Minkowski functional of its unit ball.

For any norm $\|\cdot\|$, define the dual norm $\|\cdot\|_*$ by $\|z\|_* = \sup\{\langle x, z \rangle : \|x\| \leq 1\}$. For any convex body \mathcal{K} with $0 \in \mathcal{K}$, define the polar as $\mathcal{K}^* = \{y : \sup_{x \in \mathcal{K}} \langle x, y \rangle \leq 1\}$. For any symmetric convex body, the dual norm of $\|\cdot\|_{\mathcal{K}}$ is given by $\|\cdot\|_{\mathcal{K}^*}$. The dual norm of $\|\cdot\|_2$ is itself, while the dual norm of $\|\cdot\|_\infty$ is $\|\cdot\|_1$ (and vice versa). For matrix norms, the dual of $\|\cdot\|_F$ is itself, while the dual norm of $\|\cdot\|$ is $\|\cdot\|_1$ (and vice versa).

For any random variable X , let $\mathbb{E} X$ denote its expectation (if it exists) and let $\mathbb{V}(X) = \mathbb{E} X^2 - (\mathbb{E} X)^2$ denote the variance (if it exists). For any $\mu \in \mathbb{R}^d$ and $\Sigma \geq 0$, denote by $\mathcal{N}(\mu, \Sigma)$, the normal distribution with mean μ and covariance Σ .

We also record here some notation from quantum information. A density matrix ρ is a positive semidefinite matrix with trace one i.e. $\rho \geq 0$ and $\text{tr} \rho = 1$. Measurements in quantum information are specified by POVMs which are PSD matrices A_i such that $\sum_i A_i = I$ (POVMs are more general than this, but for the context of this paper it suffices to consider this definition). For any density matrix, upon measuring ρ with respect to the POVM $\{A_i\}$, one gets outcome i with probability $\text{tr} \rho A_i$. For any density matrix, define the von Neumann entropy as $S(\rho) = -\text{tr}(\rho \log \rho)$.

3 PROOF OF MAIN THEOREM

In this section we prove the following main partial coloring theorem. Then in Section 3.1 we use it to deduce Theorem 1.3. Theorem 1.5 can then be proved by a simple modification of the proof of Theorem 1.3.

THEOREM 3.1 (MAIN PARTIAL COLORING THEOREM). *Let $A_1, \dots, A_n \in \mathbb{R}^{d \times d}$ be symmetric. There is a partial fractional coloring $x \in [-1, 1]^n$ such that $\Omega(n)$ indices i have $|x_i| = 1$ and*

$$\left\| \sum_{i \leq n} x_i A_i \right\| \leq \left\| \sum_{i \leq n} A_i^2 \right\|^{1/2} \cdot O\left(\sqrt{1 + \log \frac{\text{Tr} \sum_{i \leq n} A_i^2}{\sqrt{n} \|\sum_{i \leq n} A_i^2\|}}\right).$$

Furthermore, there is a randomized polynomial time algorithm which finds such a coloring with high probability.

It is a folklore observation⁵ that if $n \gg d^2$ one can find a partial coloring with zero discrepancy by linear programming, so Theorem 3.1 is interesting when $n \ll d^2$.

We now assemble our main tools for the proof of Theorem 3.1. Our first lemma shows that if a low-discrepancy partial fractional coloring of A_1, \dots, A_n does not exist then A_1, \dots, A_n induce a scheme to compress a large subset of $\{\pm 1\}^n$ into the $d \times d$ nuclear norm ball. The nuclear norm appears because it is dual to the norm in which we are measuring discrepancy, namely spectral norm. Using convex duality, in Section 4 we actually prove the following more general statement which applies to any norm and its dual, in hope that it is useful in future work.

⁵Thanks to Raghu Meka for making us aware of this.

LEMMA 3.2 (COMPRESS OR COLOR). *Let \mathcal{K} be a symmetric convex body in \mathbb{R}^m and let $\|\cdot\|_{\mathcal{K}}$ be its associated norm. Let $v_1, \dots, v_n \in \mathbb{R}^m$. For every $\epsilon, \Delta > 0$, either*

- *there is a partial fractional coloring $x \in [-1, 1]^n$ such that $\{|i : |x_i| = 1\}| \geq \epsilon n$ and $\|\sum_{i \leq n} x_i v_i\|_{\mathcal{K}} \leq \Delta$, and, furthermore, with probability $\Omega(1/n)$ over uniformly random choice of $g \in \{\pm 1\}^n$ such a coloring is given by optimizer of the following convex program:*

$$\max_x \langle x, g \rangle \text{ such that } \left\| \sum_{i \leq n} x_i v_i \right\|_{\mathcal{K}} \leq \Delta \text{ and } x \in [-1, 1]^n, \text{ or,}$$

- *for at least $2^n/2$ choices of $g \in \{\pm 1\}^n$ there is a vector $y_g \in \mathbb{R}^m$ with $\|y_g\|_{\mathcal{K}^*} = 1$, a set $I_g \subseteq [n]$ with $|I_g| \leq \epsilon n$, and for $i \in [n] \setminus I_g$ numbers $\Delta_{ig} \geq \Omega(\Delta)$ such that $\langle y_g, v_i \rangle = \frac{\Delta_{ig}}{n} g_i$.*

In light of Lemma 3.2, to show in the proof of Theorem 3.1 that there is a partial fractional coloring with discrepancy Δ , we can instead rule out a mapping from $g \in \{\pm 1\}^n$ to matrices Y_g with $\|Y_g\|_1 = 1$ (since the nuclear norm is dual to the spectral norm) such that $\langle Y_g, A_i \rangle \approx \frac{\Delta}{n} \cdot g_i$ for at least $(1 - \epsilon)n$ indices $i \in [n]$.

By a simple transformation of the Y_g s and A_i s, we can assume $Y_g \geq 0$, and hence that Y_g is a d -dimensional density matrix, and that $\text{Tr } A_i = 0$. We can then interpret $\{Y_g\}_{g \in \{\pm 1\}^n}$ as a strategy for Alice in a one-way quantum protocol for the n -bit index function, where Bob's measurements are the A_i 's. In our discussion of such protocols so far, we have always assumed that $\|A_i\| \leq 1$; note that we do not make this assumption here. It turns out that the weaker assumption on $\sum_{i \leq n} A_i^2$ suffices to build the repeated protocol we need to prove our communication lower bound.

Since our communication lower bounds only apply to protocols where Alice communicates a pure state, we use following lemma to round $\{Y_g\}_{g \in \{\pm 1\}^n}$ to pure states $\{y_g\}_{g \in \{\pm 1\}^n}$. The cost is that the fluctuations in this randomized rounding scheme are governed by $\text{Tr} \sum_{i \leq n} A_i^2 = \sum_{i \leq n} \|A_i\|_F^2$ in addition to $\|\sum_{i \leq n} A_i^2\|$. The assumption $\|A_i\|_F \leq n^{1/4}$ is needed to control the first term.

LEMMA 3.3 (PURIFY THEN SKETCH). *Let $n, d, \delta > 0$. Let $A_1, \dots, A_n \in \mathbb{R}^{d \times d}$ have $\text{Tr } A_i = 0$ and $A = \sum_{i=1}^n A_i^2$. Let $\{Y_g\}_{g \in \{\pm 1\}^n}$ be density matrices.*

For every integer $r > 0$, there exist symmetric matrices $B_1, \dots, B_n \in \mathbb{R}^{r \times r}$ such that

$$\left\| \sum_{i=1}^n B_i^2 \right\| \leq O_{\delta} \left(\|A\| + \frac{\min(\sqrt{n}, d) \cdot \text{Tr } A}{r} \right)$$

and such that for at least $\frac{3}{4} \cdot 2^n$ of $g \in \{\pm 1\}^n$ there exists an r -dimensional pure state y_g (i.e. a vector $y_g \in \mathbb{R}^r$ with $\|y_g\| = 1$) and a number $c_g \geq \Omega_{\delta}(1)$ such that for at least $(1 - \delta)n$ indices $i \in [n]$,

$$\left| \langle y_g y_g^\top, B_i \rangle - c_g \langle Y_g, A_i \rangle \right| \leq O_{\delta} \left(\frac{\|A\|}{nr} + \frac{\min(\sqrt{n}, d) \text{Tr } A}{nr^2} \right)^{1/2}.$$

Lastly, we prove the following lemma using tools from quantum information and communication complexity – see Section 5 for a more thorough discussion.

LEMMA 3.4. *There is a universal constant $\delta > 0$ such that the following holds for all integers $n, m > 0$. Let $\mathcal{S} \subseteq \{\pm 1\}^n$ have size at least $2^{\delta n}$. Suppose that $\{y_g\}_{g \in \mathcal{S}}$ is a collection of m -qubit pure states*

such that for some symmetric matrices A_1, \dots, A_n , subsets $I_g \in \binom{[n]}{\delta n}$ and numbers $\{\eta_{ig} > 0\}_{g \in \mathcal{S}, i \in [n] \setminus I_g}$, for all $i \in [n] \setminus I_g$ it holds that $\langle y_g y_g^\top, A_i \rangle = \eta_{ig} \cdot g_i$. Let

$$\eta^2 = \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} \eta_{ig}^2}{\|\mathbb{E}_{i \sim [n]} A_i^2\|}.$$

If $\eta \geq C/\sqrt{n}$ for some universal $C > 0$, then

$$m \geq \log \frac{1}{\eta^2} + \Omega(\eta^2 \cdot n).$$

With our tools in hand, we can prove Theorem 3.1.

PROOF OF THEOREM 3.1. Let

Good Compression(Δ): for $2^n/2$ choices of $g \in \{\pm 1\}^n$ there is a matrix Y_g with $\|Y_g\|_1 = 1$, a set $I_g \subseteq [n]$ with $|I_g| \leq \epsilon n$, and numbers $\{\Delta_{ig}\}_{i \in [n] \setminus I_g}$ with $\Delta_{ig} \geq \Delta$ such that $\langle Y_g, A_i \rangle = \frac{\Delta_{ig}}{n} g_i$ for $i \in [n] \setminus I_g$

By the Compress or Color Lemma 3.2, it is enough to show that for a small-enough constant $\epsilon > 0$, if **Good Compression(Δ)** occurs, then

$$\Delta \leq \left\| \sum_{i \leq n} A_i^2 \right\|^{1/2} \cdot O \left(\min \left\{ \sqrt{1 + \log \frac{\text{Tr} \sum_{i \leq n} A_i^2}{\sqrt{n} \|\sum_{i \leq n} A_i^2\|}}, \sqrt{n \cdot 2^{-\Omega(n/d^2)}} \right\} \right). \quad (3.1)$$

In that case, if Δ is larger than in (3.1), then with probability $\Omega(1/n)$ over choice of $g \in \{\pm 1\}^n$, the semidefinite program

$$\max_{x \in [-1, 1]^n} \langle x, g \rangle \text{ such that } \left\| \sum_{i \leq n} x_i A_i \right\| \leq \Delta$$

finds a fractional partial coloring with ϵn integer entries and discrepancy at most Δ .

Suppose **Good Compression(Δ)** occurs, for some $\Delta > 0$. Let $Y_g^+ \geq 0$ be the positive semidefinite part of Y_g and $Y_g^- \leq 0$ the negative definite part, so that $Y_g = Y_g^+ + Y_g^-$. By replacing Y_g and A_i with the following block matrices:

$$Y_g \rightarrow \begin{pmatrix} Y_g^+ & 0 \\ 0 & -Y_g^- \end{pmatrix} \text{ and } A_i \rightarrow \begin{pmatrix} A_i & 0 \\ 0 & -A_i \end{pmatrix},$$

(and replacing d with $2d$) we may assume that $Y_g \geq 0$ with $\text{Tr } Y_g = 1$ and $\text{Tr } A_i = 0$. Note that $\langle Y_g, A_i \rangle$ and $\|\sum_{i \leq n} A_i^2\|$ are preserved by this transformation, and $\text{Tr} \sum_{i \leq n} A_i^2$ grows by a factor of 2.

Let $A = \sum_{i \leq n} A_i^2$. By the Purify-then-Sketch lemma 3.3, for any choice of integer $r > 0$ and any $\delta > 0$ there are symmetric matrices $B_1, \dots, B_n \in \mathbb{R}^{r \times r}$ such that

$$\left\| \sum_{i \leq n} B_i^2 \right\| \leq O_{\delta} \left(\|A\| + \frac{\min(\sqrt{n}, d) \text{Tr } A}{r} \right)$$

and for at least $\frac{3}{4} \cdot 2^n$ choices of $g \in \{\pm 1\}^n$ there is a pure state $y_g \in \mathbb{R}^r$ and a number $c_g \geq \Omega_{\delta}(1)$ such that for at least $(1 - \delta)n$ indices $i \in [n]$,

$$\left| \langle y_g y_g^\top, B_i \rangle - c_g \langle Y_g, A_i \rangle \right| \leq O_{\delta} \left(\frac{\|A\|}{nr} + \frac{\min(\sqrt{n}, d) \text{Tr } A}{nr^2} \right)^{1/2}.$$

Now, if

$$\Delta^2 \geq O_\delta(1) \cdot \|A\| \cdot \max\left(\frac{n}{r}, \frac{n}{r} \cdot \frac{\min(\sqrt{n}, d) \operatorname{Tr} A}{r}\right), \quad (3.2)$$

then for at least $2^n/4$ choices of g there is a subset $I_g \subseteq [n]$ with $|I_g| \leq (\varepsilon + \delta)n$ such that for $i \in [n] \setminus I_g$,

$$\langle y_g y_g^\top, B_i \rangle = \frac{\Delta_{ig}}{n} \cdot g_i$$

where $\Delta_{ig} \geq \Omega(\Delta)$. Let us call this set of g 's \mathcal{S} .

Let

$$\eta^2 = \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} (\Delta_{ig}/n)^2}{\|\mathbb{E}_{i \sim [n]} B_i^2\|}.$$

Henceforth taking δ to be a small-enough universal constant, by Lemma 3.4, either $\eta \leq O(1/\sqrt{n})$ or $r \geq 2^{\Omega(\eta^2 n)} \cdot \frac{1}{\eta^2}$. We treat the two cases separately.

Case 1A: $\eta \leq O(1/\sqrt{n})$ Using $\Delta_{ig} \geq \Omega(\Delta)$, in this case we have

$$\Delta^2 \leq \left\| \mathbb{E}_{i \sim [n]} B_i^2 \right\| \leq O\left(\|A\| + \frac{\min(\sqrt{n}, d) \operatorname{Tr} A}{r}\right)$$

Case 1B: $r \geq 2^{\Omega(\eta^2 n)} \cdot \frac{1}{\eta^2}$ Using the definition of η , we have

$$\log\left(r \cdot \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} (\Delta_{ig}/n)^2}{\|\mathbb{E}_{i \sim [n]} B_i^2\|}\right) \geq \Omega(n) \cdot \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} (\Delta_{ig}/n)^2}{\|\mathbb{E}_{i \sim [n]} B_i^2\|}$$

This rearranges to

$$\log\left(\frac{r}{n} \cdot \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} \Delta_{ig}^2}{\|\sum_{i \in n} B_i^2\|}\right) \geq \Omega(1) \cdot \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} \Delta_{ig}^2}{\|\sum_{i \in n} B_i^2\|}$$

which gives

$$\log \frac{r}{n} \geq \Omega(1) \cdot \frac{\mathbb{E}_{g \sim \mathcal{S}, i \sim [n] \setminus I_g} \Delta_{ig}^2}{\|\sum_{i \in n} B_i^2\|},$$

so, rearranging and using our bound on $\|\sum_{i \in n} B_i^2\|$, and that $\Delta_{ig} \geq \Omega(\Delta)$, we get

$$\Delta^2 \leq O\left(\|A\| + \frac{\min(\sqrt{n}, d) \operatorname{Tr} A}{r}\right) \cdot \log \frac{r}{n}.$$

Now let us choose

$$r = \max\left(n, \frac{\min(\sqrt{n}, d) \operatorname{Tr} A}{\|A\|}\right),$$

so that putting together (3.2) with cases 1A and 1B, we find

$$\Delta^2 \leq O(\|A\|) \max\left(1, \log \frac{\min(\sqrt{n}, d) \operatorname{Tr} A}{n \|A\|}\right). \quad \square$$

3.1 Matrix Spencer for Moderate-Rank Matrices

In this section we use Theorem 3.1 to prove the following theorem.

THEOREM 3.5. *Let $d \geq n$ and $A_1, \dots, A_n \in \mathbb{R}^{d \times d}$ be symmetric matrices with $\|A_i\| \leq 1$ and $\|A_i\|_F^2 \leq \sqrt{d}$ for all $i \in [n]$. There is a coloring $x \in \{\pm 1\}^n$ such that*

$$\left\| \sum_{i \in n} x_i A_i \right\| \leq O\left(\sqrt{n \log\left(\frac{d}{n}\right)}\right).$$

Furthermore, there is a (randomized) polynomial time algorithm which finds such a coloring with high probability.

PROOF. We will get a full coloring for this setting by iteratively applying Lemma 3.1. Consider the first round of partial coloring. Note that since $\|A_i\| \leq 1$, we have $\|\sum_{i \in n} A_i^2\| \leq n$. From Lemma 3.1, we get that there is a partial coloring x with cn co-ordinates such that $x_i \in \{-1, 1\}$ and

$$\left\| \sum_{i \in n} x_i A_i \right\| \leq \left\| \sum_{i \in n} A_i^2 \right\|^{1/2} \cdot O\left(\sqrt{1 + \log \frac{\operatorname{Tr} \sum_{i \in n} A_i^2}{\sqrt{n} \|\sum_{i \in n} A_i^2\|}}\right).$$

By hypothesis on $\|A_i\|_F^2 = \operatorname{Tr} A_i^2$, we have $\operatorname{Tr} \sum_{i \in n} A_i^2 \leq n \cdot d^{1/2}$.

$$\left\| \sum_{i \in n} x_i A_i \right\| \leq O\left(\left\| \sum_{i \in n} A_i^2 \right\|^{1/2} \cdot \sqrt{\log \frac{2\sqrt{dn}}{\|\sum_{i \in n} A_i^2\|}}\right).$$

Since $x \log(2\sqrt{dn}/x) \leq O(n \log(d/n))$ for $0 \leq x \leq n$, we get a partial coloring with discrepancy $O(\sqrt{n \log(d/n)})$.

Given a partial coloring x , we move to the next round of partial coloring by replacing A_i by

$$A'_i = \operatorname{sign}(x_i) \cdot (1 - |x_i|) A_i.$$

We ignore the co-ordinates corresponding to zero matrices. Since $(1 - |x_i|) \leq 1$, the new matrices still satisfy the requirements on the spectral norm and Frobenius norm. Furthermore, since cn co-ordinates are integral, the number of matrices is now $(1 - c)n$. Thus, we get a partial coloring with y such that

$$\left\| \sum_{i \in n} y_i A'_i \right\| \leq \left\| \sum_i A'_i \right\|^{1/2} O\left(1 + \log \frac{\operatorname{Tr} \sum_i A'_i}{\sqrt{(1 - c)n} \|\sum_i A'_i\|}\right).$$

Arguing as before, noting that $\operatorname{Tr} \sum_i A'_i \leq (1 - c)n\sqrt{d}$, we get

$$\left\| \sum_{i \in n} y_i A'_i \right\| \leq O\left(\sqrt{(1 - c)n \log\left(\frac{d}{(1 - c)n}\right)}\right).$$

Then, consider the partial coloring with co-ordinates $z_i = x_i + \operatorname{sign}(x_i)(1 - |x_i|) y_i$. First note that we have $c + (1 - c)c/2$ co-ordinates that are integral. To see this note that for half the integral co-ordinates in y , we must have $\operatorname{sign}(y_i) = 1$ (replacing y with $-y$ if necessary). For such y_i , we have $x_i + \operatorname{sign}(x_i)(1 - |x_i|) y_i = \operatorname{sign}(x_i)$ which is integral. Furthermore,

$$\begin{aligned} \left\| \sum_i z_i A_i \right\| &\leq \left\| \sum_i x_i A_i \right\| + \left\| \sum_i y_i A'_i \right\| \\ &\leq O\left(\sqrt{n \log\left(\frac{d}{n}\right)} + \sqrt{(1 - c)n \log\left(\frac{d}{(1 - c)n}\right)}\right). \end{aligned}$$

Iterating this inductively, we get that the discrepancy is bounded by

$$O\left(\sqrt{n} \sum_i (1 - c)^{i/2} \sqrt{\log\left(\frac{d}{(1 - c)^i n}\right)}\right) = O\left(\sqrt{n \log\left(\frac{d}{n}\right)}\right)$$

as required. \square

4 COMPRESS OR COLOR

In this section we prove Lemma 3.2. For this, we would like to understand the optimizers of the program given by $\max_x \langle x, g \rangle$ for x in the set of low discrepancy vectors for a set of vectors v_1, \dots, v_n under some norm and $g \in \{-1, 1\}^n$. Lemma 3.2 follows from the following proposition which connects optimizers to such programs with few integral coordinates to unit vectors under the dual norm whose inner products with v_i correlate well with the coordinates of g .

LEMMA 4.1. *Let $v_1, \dots, v_n \in \mathbb{R}^m$ and $\|\cdot\|$ be a norm. Let $\varepsilon \in (0, 1)$ and $\Delta > 0$. Then, for any $g \in \{-1, 1\}^n$ such that any optimizer to the program*

$$\begin{aligned} \max_{x \in [-1,1]^n} \langle x, g \rangle \\ \left\| \sum_i x_i v_i \right\| \leq \Delta \end{aligned}$$

has less than εn coordinates in $\{-1, 1\}$, there exists a y_g with $\|y_g\|_ = 1$ and $I_g \subseteq [n]$ with $|I_g| \leq \varepsilon n$ such that*

$$\langle y_g, v_i \rangle g_i \geq \frac{\Delta}{n}$$

for all $i \notin I_g$.

PROOF. Let $g \in \{-1, 1\}^n$ as in the statement of the lemma i.e. such that any optimizer to the program

$$\begin{aligned} \max_{x \in [-1,1]^n} \langle x, g \rangle \\ \left\| \sum_i x_i v_i \right\| \leq \Delta \end{aligned}$$

has less than εn coordinates in $\{-1, 1\}$. Let x_g such an optimizer. Denote by $T(x)$ the set $\{i \in [n] : |x_i| = 1\}$. Note that we have $|T(x)| \leq \varepsilon n$. Denote by $K_\Delta = \{x : \|\sum_i x_i v_i\| \leq \Delta\}$ and $[-1, 1]^{T(x)} = \{x : |x_i| \leq 1 \quad \forall i \in T(x)\}$ and $Q = K_\Delta \cap [-1, 1]^{T(x)}$. Note that the following convex program

$$\max_{x \in Q} \langle x, g \rangle$$

has the same optimizer and value as the original program. This is due to the fact that removing non-tight constraints in a convex program does not change the optimum and optimizer.

Further, observe that $\langle x_g, g \rangle \leq n$ since $x_g, g \in [-1, 1]^n$. Thus, $\max_{x \in Q} \langle x, g \rangle \leq n$, which implies that $g/n \in Q^*$, the polar body of Q . Thus, we have $g/n = \lambda y_1 + (1 - \lambda) y_2$ with $\lambda \in (0, 1)$, and $y_1 \in K_\Delta^*$ and $y_2 \in B_1^{T(x)}$. Here $B_1^{T(x)} = \{y : y_i = 0 \quad \forall y_i \notin T(x), \sum_i |y_i| = 1\}$ which is the polar body of $[-1, 1]^{T(x)}$. Since y_2 has only zero coordinates for $i \notin T(x)$, we have $g_i/n = (\lambda y_2)_i$. But note that $K_\Delta^* = \Delta^{-1} \{(\langle v_1, y \rangle, \dots, \langle v_n, y \rangle) : \|y\|_* \leq 1\}$. Thus, we get the required result. \square

5 ONE-WAY COMMUNICATION IN THE SMALL-ADVANTAGE REGIME

In this section we study (quantum) communication complexity of the index function. We adopt terminology from the quantum information literature, where a one-way protocol for the index function is called a (quantum) random access code.

Definition 5.1 (Random Access Code). A (m, n, ε) random access code is a map from messages $g \in \{\pm 1\}^n$ to distributions over m -bit codewords y_g together with a family of decoding procedures $D_1, \dots, D_n : \{\pm 1\}^m \rightarrow \{\pm 1\}$ such that for all g and all i , the decoding procedure D_i run on input y_g outputs g_i with probability at least $1 - \varepsilon$ (over the random choice of encoding of g).

Definition 5.2 (Quantum random access code). A *quantum* random access code is a map from message g to quantum states $\{y_g\}_{g \in \{\pm 1\}^n}$ together with a family of quantum decoding procedures $D_1, \dots, D_n : \{\pm 1\}^m \rightarrow \{\pm 1\}$. We call the code *pure* if y_g is a pure state; otherwise it may be a mixed state. The decoding procedure D_i have the property that for all g and all i , D_i run on input y_g outputs g_i with probability at least $1 - \varepsilon$, where the probability is over any classical randomness in the possibly mixed state y_g as well as the randomness in the outcomes of measurements made by D_i .

In the regime that ε is a constant independent of n, m , classical information theory shows that $m \geq \Omega(n)$ for classical codes. This is not so obvious for the quantum case, since the decoding procedure D_i may destroy the state y_g and make other bits unreadable. Nonetheless, a clever application of the Holevo bound together with an inductive argument shows that $m \geq \Omega(n)$ in quantum case as well [2].

For applications to discrepancy, we are interested in the case that ε is close to $1/2$ – i.e. $\varepsilon = 1/2 - \eta$ for some $\eta = \eta(n) \rightarrow 0$. Of particular interest is the regime $\eta = \Theta(1/\sqrt{n})$ – this is the most relevant regime for Spencer-style discrepancy bounds, and it also turns out to be interesting from the perspective of random access codes, whose behavior is rather different for $\eta \gg 1/\sqrt{n}$ and $\eta \ll 1/\sqrt{n}$.

To carry out our application to discrepancy, we need lower bounds for somewhat weaker notions of random access codes.

Definition 5.3 (Weakness). A (m, n, ε) random access code is called (δ, t) -weak (for a subset $\mathcal{S} \subseteq \{\pm 1\}^n$) if it contains messages y_g only for $g \in \mathcal{S}$, where $\log |\mathcal{S}| \geq n - t$, and if for each g at least a $(1 - \delta)$ fraction of the decoding procedures D_i yield the bit g_i with probability $1 - \varepsilon$. (So for each g there may be as many as δn bad coordinates which are not decoded by their corresponding decoding procedures.)

We prove the following main result for both classical and pure quantum random access codes. Adapting the same proof, afterwards we prove Lemma 3.4.

THEOREM 5.4 (LOWER BOUND FOR LOW-SIGNAL (PURE) RANDOM ACCESS CODES). *Let $\{y_g\}_{g \in \{\pm 1\}^n}$ be an $(m, n, \frac{1}{2} - \eta)$ classical random access code or pure quantum random access code. Then*

$$m \geq \log n - \log \log \frac{1}{\eta^2} - O(1).$$

Furthermore, if $\eta \geq 10/\sqrt{n}$,

$$m \geq \log \frac{1}{\eta^2} + \Omega(\eta^2 n).$$

Furthermore, the same inequalities hold if $\{y_g\}$ is (δ, t) -weak, so long as $\delta \leq \delta_0$ and $t \leq t_0 n$ for some universal constants δ_0, t_0 .

We conjecture that the same bounds as in Theorem 5.4 hold also for general quantum random access codes, where y_g may be a mixed state – such bounds would imply the Matrix Spencer conjecture.

The proof of Theorem 5.4 has two parts. First, we generalize the argument of [2] to the case of weak quantum random access codes (requiring only a slight adaptation of the arguments of [2]). Then, we prove Theorem 5.4 by applying the resulting lemma not to the random access code we start with, but instead to the code we get by amplifying the success probability by repeating the code.

LEMMA 5.5 (LOWER BOUND FOR WQRACs, ADAPTED FROM [2]). *If there exists an (m, δ, ϵ) -WQRAC for a subset $\mathcal{S} \subset \{\pm 1\}^n$ with $\log |\mathcal{S}| = n - t$, then*

$$m \geq (1 - H(\delta) - H(\epsilon)) \cdot n - t.$$

Now we prove Theorem 5.4 and its refinement Lemma 3.4, deferring the proof of Lemma 5.5 to the end of this section.

PROOF OF THEOREM 5.4. Let $r > 0$ be an integer. Consider the *amplified* random access code which uses

- r independent draws y_g^1, \dots, y_g^r from the distribution of messages encoding g , in the classical case, and
- $y_g^{\otimes r}$, that is, r copies of the state y_g , in the quantum case,

as the encoding of g . To decode the bit i , run the decoding procedure D_i on each copy and take a majority vote. This gives a (δ, t) -weak random access code for \mathcal{S} with failure probability $\epsilon' = \mathbb{P}(\text{Bin}(r, 1/2 + \eta) < r/2) \leq \exp(-r\eta^2)$. Take $r = 1/\eta^2$, so that this is at most $1/e$.

We claim that in classical case, the message y_g^1, \dots, y_g^r can be expressed using just

$$\log \left(\frac{2^m + 1/\eta^2}{1/\eta^2} \right) \leq \min \left(2^m, \frac{1}{\eta^2} \right) \cdot \log \left(e \cdot \frac{2^m + \frac{1}{\eta^2}}{\min(2^m, \frac{1}{\eta^2})} \right) \quad (5.1)$$

bits, and similarly with at most the above number of qubits for the pure quantum case.

Classical: The majority-vote decoding procedure only needs to know the frequency of each of the 2^m possible messages among $y_g^1, \dots, y_g^{1/\eta^2}$. By a “stars and bars” argument, the number of such frequency-counts is at most $\binom{2^m + 1/\eta^2}{1/\eta^2}$.

Quantum: The state $y_g^{\otimes 1/\eta^2}$ lies in the *symmetric subspace*, $\text{span}\{x^{\otimes 1/\eta^2} : x \in \mathbb{R}^{2^m}\}$. This subspace has dimension at most $\binom{2^m + 1/\eta^2}{1/\eta^2}$ [19] (by the same stars and bars argument).

Applying Lemma 5.5 to this amplified code gives us the following bound.

$$\min \left(2^m, \frac{1}{\eta^2} \right) \cdot \log \left(e \cdot \frac{2^m + \frac{1}{\eta^2}}{\min(2^m, \frac{1}{\eta^2})} \right) \quad (5.2)$$

$$\geq (1 - H(\delta) - H(1/e))n - t \quad (5.3)$$

$$\geq \frac{n}{2}. \quad (5.4)$$

where we have used the hypotheses on δ and t .

Suppose first that we take the bound

$$2^m \log \left(e \cdot \frac{2^m + 1/\eta^2}{2^m} \right) \geq n/2.$$

Then we have

$$2^m \cdot \left(\log \frac{1}{\eta^2} + O(1) \right) \geq \frac{n}{2}.$$

Hence,

$$m \geq \log n - \log \log \frac{1}{\eta^2} - O(1)$$

as desired.

Now, suppose that $\eta \geq 10/\sqrt{n}$. Let $c = \log 1/\eta^2 - m$, so that $2^m = 2^{-c}/\eta^2$. Suppose $c \geq 0$, so that $2^m \leq 1/\eta^2$. Then we have

$$2^m \left(\log \frac{1}{\eta^2 2^m} + 3 \right) \geq \frac{n}{2}.$$

By definition of c , we get $2^m(c + 3) \geq n/2$, which gives $m \geq \log n + \log(1/(c + 3)) - 1$. On the other hand, by hypothesis $2^m = 2^{-c}/\eta^2$ so $m \leq \log(1/\eta^2) - c$. By hypothesis on η , this means $m \leq \log n - \log(100) - c$. This is a contradiction for $c \geq 0$; we must have $c < 0$.

Put differently, we may assume that $1/\eta^2 \leq 2^m$. Using (5.4) again, we obtain

$$\frac{1}{\eta^2} \log \left(e \cdot \frac{2^m + \frac{1}{\eta^2}}{\frac{1}{\eta^2}} \right) \geq n/2.$$

This gives

$$\frac{2^m + \frac{1}{\eta^2}}{\frac{1}{\eta^2}} \geq 2^{\eta^2 n/2 - 2}$$

and, rearranging:

$$2^m \geq \frac{1}{\eta^2} \cdot 2^{\eta^2 n/2 - 2} - \frac{1}{\eta^2} = \frac{1}{\eta^2} \cdot \left(2^{\eta^2 n/2 - 2} - 1 \right)$$

Taking logs,

$$m \geq \log \frac{1}{\eta^2} + \log \left(2^{\eta^2 n/2 - 2} - 1 \right) \geq \log \frac{1}{\eta^2} + \Omega(\eta^2 n)$$

where we used the assumption $\eta \geq 10/\sqrt{n}$ for the second inequality. \square

PROOF OF LEMMA 3.4. The proof is identical to that of Theorem 5.4, except that we construct the decoding procedures for the amplified code as follows.

Decoding the amplified code: For each i , the matrix A_i induces the following measurement procedure: measure in the eigenbasis a_1, \dots, a_n of A_i , and on receiving outcome j , output the eigenvalue of A_i associated to a_j . Given r copies of the state y_g , to decode the i -th bit, run the aforementioned measurement procedure on each copy of y_g and average the results. Output 1 if the sum is positive and -1 otherwise.

Analysis of decoding: We claim that for some choice of $r = O(\|\mathbb{E}_{i \sim [n]} A_i^2\|/\mathbb{E}_g \eta_g^2)$, the above decoding procedure yields a weak random access code for a subset $\mathcal{S}' \subseteq \mathcal{S}$ of size $|\mathcal{S}'| \geq (1 - \delta)|\mathcal{S}|$, with failure probability at most $1/e$. After this is established, the proof can proceed as in Theorem 5.4, with $\eta = \Theta(\sqrt{\mathbb{E}_g \eta_g^2}/\|\mathbb{E}_{i \sim [n]} A_i^2\|)$.

In decoding the i -th bit, the decoding procedure produces a sum of i.i.d. random variables X_1, \dots, X_r , each with mean $\eta_g a_i$ and

variance $\langle y_g y_g^\top, A_i^2 \rangle$. The average $X = \sum_{i=1}^r X_i/r$ has $\mathbb{E} X = \eta_g g_i$ and variance $\mathbb{V}(X) = \langle y_g y_g^\top, A_i^2 \rangle/r$. By Chebyshev's inequality, for fixed i and g , the probability of incorrectly decoding the bit g_i is at most

$$\mathbb{P}(|X - \mathbb{E} X| > \eta_g) \leq \frac{\eta_g^2}{\langle y_g y_g^\top, A_i^2 \rangle/r}.$$

This is at most $1/e$ so long as $r \geq e \cdot \langle y_g y_g^\top, A_i^2 \rangle / \eta_g^2$. By Markov's inequality, for any $\delta > 0$, if we choose $r \geq C(\delta) \|\mathbb{E}_{i \sim [n]} A_i^2\| / \mathbb{E}_g \eta_g^2$, we will have $r \geq e \cdot \langle y_g y_g^\top, A_i^2 \rangle / \eta_g^2$ for at least a $(1 - \delta)$ fraction of pairs (g, i) for which the assumptions of the theorem apply. \square

5.1 Proof of Lemma 5.5

To prove Lemma 5.5, we use the entropy coalescence lemma of [2], which is itself a corollary of the Holevo bound.

LEMMA 5.6 (ENTROPY COALESCENCE LEMMA, [2]). *Let ρ_{-1} and ρ_1 be two density matrices and let $\rho = (1 - \beta) \cdot \rho_{-1} + \beta \cdot \rho_1$ be their mixture for some $\beta \in (0, 1)$. If there is a measurement with outcome -1 or 1 such that making the measurement on ρ_b yields the bit b with probability at least p , then*

$$S(\rho) \geq (1 - \beta) \cdot S(\rho_{-1}) + \beta \cdot S(\rho_1) + (H(\beta) - H(p))$$

where H is the binary entropy function and S is the von Neumann entropy.

Now we can prove Lemma 5.5 essentially by following the argument of [2] while throwing out the ε -fraction of g 's where Y_g is not decodable to g .

PROOF OF LEMMA 5.5. For each $g \in \mathcal{S}$, there is a set of at most $\delta \cdot n$ bad coordinates $i \in [n]$ for which the corresponding decoding procedures D_i do not produce the bit g_i with probability $1 - \varepsilon$. Let $I_g \in \binom{[n]}{\delta n}$ denote this subset. Let $I \in \binom{[n]}{\delta n}$ denote the subset that is the bad set for the largest number of strings $g \in \mathcal{S}$, and let

$$S_I = \{g \in \mathcal{S} \mid I_g = I\}$$

Clearly, $|S_I| \geq |\mathcal{S}| / \binom{n}{\delta n}$.

Without loss of generality, let us assume that the set I consists of the last δn coordinates, i.e., $I = \{(1 - \delta)n, \dots, n\}$. Let Θ denote the uniform distribution over $g \in \mathcal{S}_I$. Let $g \in \{\pm 1\}^n$ denote a random sample from the distribution Θ . For every $L < (1 - \delta) \cdot n$, let Θ_L denote the marginal distribution of Θ over the first L coordinates. For every $\alpha \in \{\pm 1\}^L$, let us define

$$\rho_\alpha = \mathbb{E}_{g \sim \Theta} [Y_g \mid g_L = \alpha]$$

By definition, we can write

$$\rho_\alpha = \mathbb{P}[g_{L+1} = +1 \mid g_L = \alpha] \cdot \rho_{\alpha,+1} + \mathbb{P}[g_{L+1} = -1 \mid g_L = \alpha] \cdot \rho_{\alpha,-1}$$

Applying the entropy coalescence lemma Lemma 5.6, we get that

$$\begin{aligned} S(\rho_\alpha) &\geq \mathbb{P}[g_{L+1} = +1 \mid g_L = \alpha] \cdot S(\rho_{\alpha,+1}) + \mathbb{P}[g_{L+1} = -1 \mid g_L = \alpha] \cdot S(\rho_{\alpha,-1}) \\ &\quad + H(\mathbb{P}[g_{L+1} = +1 \mid g_L = \alpha]) - H(\varepsilon) \end{aligned}$$

Averaging the above inequality over α drawn from the distribution Θ_L ,

$$\mathbb{E}_{\alpha \sim \Theta_L} [S(\rho_\alpha)] \geq \mathbb{E}_{\alpha' \sim \Theta_{L+1}} [S(\rho_{\alpha'})] + H(g_{L+1} \mid g_L) - H(\varepsilon)$$

Summing up the above inequality over $L = 0, \dots, (1 - \delta)n - 1$,

$$\begin{aligned} S(\rho_\emptyset) &\geq \mathbb{E}_{\alpha \sim \Theta} [S(\rho_\alpha)] + \sum_{L=1}^n H(g_{L+1} \mid g_L) - (1 - \delta)n \cdot H(\varepsilon) \\ &\geq \mathbb{E}_{\alpha \sim \Theta} [S(\rho_\alpha)] + H(g) - (1 - \delta)n \cdot H(\varepsilon) \\ &\geq 0 + H(g) - (1 - \delta)n \cdot H(\varepsilon) \\ &\geq \log |\mathcal{S}_I| - (1 - \delta)n \cdot H(\varepsilon) \end{aligned}$$

The lower bound follows by observing that $S(\rho_\emptyset) \leq m$ and $\binom{n}{\delta n} \leq 2^{H(\delta) \cdot n}$. \square

6 SKETCHING

In this section, we will present a random sketch that preserves evaluations of a set of quadratic forms. Fix a set of symmetric matrices $A_1, \dots, A_n \in \mathbb{R}^{D \times D}$ and a point $y \in \mathbb{R}^D$. Consider the linear sketch that samples a random matrix $S \in \mathbb{R}^{d \times D}$ with entries in $\mathcal{N}(0, 1/d)$ and maps,

$$A_i \in \mathbb{R}^{D \times D} \rightarrow S A_i S^\top \in \mathbb{R}^{d \times d}$$

and

$$y \in \mathbb{R}^D \rightarrow S y \in \mathbb{R}^d$$

We will show that this sketch approximately preserves the value of the quadratic forms $y^\top A_i y$ with good probability. The rest of this section is devoted to bounds on the expectation, variance of the sketched quadratic forms and the spectral norm of sketched matrices. These guarantees are captured by the following lemma.

LEMMA 6.1 (MAIN SKETCHING LEMMA). *Let $A_1, \dots, A_n \in \mathbb{R}^{D \times D}$ be symmetric with $\text{Tr } A_i = 0$. Let $A = \sum_{i=1}^n A_i^2$. Let $y \in \mathbb{R}^D$ be a unit vector. Finally, let $d > 0$ and let $S \in \mathbb{R}^{d \times D}$ have iid entries from $\mathcal{N}(0, 1/d)$. Then the following all hold:*

(1) **Expectation:** For all i , $\mathbb{E} \langle (S y) (S y)^\top, S A_i S^\top \rangle = \left(1 + \frac{1}{d}\right) \langle y y^\top, A_i \rangle$,

(2) **Variance:** The average variance across $i = 1, \dots, n$ is bounded:

$$\sum_{i=1}^n \mathbb{V} [\langle (S y) (S y)^\top, S A_i S^\top \rangle] \leq O\left(\frac{\|A\|}{d} + \frac{\text{Tr } A}{d^2}\right), \text{ and}$$

(3) **Spectral norm:** The following matrix has bounded spectral norm:

$$\mathbb{E} \left\| \sum_{i=1}^n (S A_i S^\top)^2 \right\| \leq O\left(\|A\| + \frac{\text{Tr } A}{d}\right).$$

The proof of the main sketching lemma may be found in the full version of the paper.

6.1 Purify then Sketch

With Lemma 6.1 in hand we can prove Lemma 3.3. We will need the following fact about the rank of solutions to semidefinite programs.

(4) THEOREM 6.2 (BARVINOK [9], PATAKI [27]). *Any compact spectahedron $\{Y : \langle Y, A_1 \rangle = b_1, \dots, \langle Y, A_m \rangle = b_m, Y \geq 0\}$ contains Y such that $\text{rank } Y \leq 4\sqrt{m}$.*

PROOF OF LEMMA 3.3. Without loss of generality, by Theorem 6.2, we may assume that $t = \max_g \text{rank } Y_g \leq \min(6\sqrt{n}, d)$. To produce y_g , we use the following algorithm:

- (1) *Purify*: Let $y'_g \in \mathbb{R}^d \otimes \mathbb{R}^t$ be a purification of Y_g . Let $A'_i = A_i \otimes \text{Id}_{t \times t}$, so that $(y'_g)^\top A'_i y'_g = \langle Y_g, A_i \rangle$.
- (2) *Sketch*: Let $S \in \mathbb{R}^{r \times dt}$ be a sketching matrix with iid entries from $\mathcal{N}(0, 1/r)$. Let $y_g = Sy'_g / \|Sy'_g\|$ and let $B_i = \frac{d}{d+1} SA'_i S^\top$.

Now we apply the main sketching lemma 6.1. Noting that $\sum_{i \leq n} (A'_i)^2 = \sum_{i \leq n} A_i^2 \otimes \text{Id}_{t \times t} = A \otimes \text{Id}_{t \times t}$, this gives for each $g \in \{\pm 1\}^n$ and each $i \leq n$,

$$\mathbb{E} \sum_{i=1}^n \left(\langle (Sy'_g)(Sy'_g)^\top, B_i \rangle - \langle Y_g, A_i \rangle \right)^2 \leq O \left(\frac{\|A\|}{r} + \frac{t \text{Tr } A}{r^2} \right). \quad (6.1)$$

For some constant C we will choose shortly, let us call g *good* if there are at least $(1 - \delta)n$ indices $i \in [n]$ such that

$$\left| \langle (Sy'_g)(Sy'_g)^\top, B_i \rangle - \langle Y_g, A_i \rangle \right| \leq C \cdot \left(\frac{\|A\|}{nr} + \frac{t \text{Tr } A}{nr^2} \right)^{1/2},$$

and, additionally, $(1/C) \leq \|Sy'_g\|^2 \leq C$. By (6.1), there is $C = C(\delta)$ such that for each g we have $\mathbb{P}(g \text{ is good}) \geq 3/4$. Therefore, $\mathbb{E}_S \mathbb{E}_{g \sim \{\pm 1\}^n} \mathbb{I}(g \text{ is good}) \geq 3/4$, and hence there is a choice of S such that $\frac{3}{4} \cdot 2^n$ g 's are good. We can obtain the pure state y_g as $Sy'_g / \|Sy'_g\|$.

The remaining claim then follows by Markov's inequality applied to $\| \sum_{i \leq n} B_i^2 \|$ (using the bound on $\mathbb{E} \| \sum_{i \leq n} B_i^2 \|$ in Lemma 6.1) and a union bound. \square

ACKNOWLEDGEMENTS

AS would like to thank Robert Kleinberg and Ayush Sekhari for enlightening conversations. We thank Tselil Schramm, Boaz Barak, Umesh Vazirani, Luca Trevisan, and Raghu Meka for several enlightening conversations as this manuscript was being prepared. We also thank an anonymous reviewer for suggesting a simpler proof of Lemma 3.2. SBH was supported by a UC Berkeley Miller Fellowship and a Simons Postdoctoral Fellowship. PR was supported by NSF grant AF 2007676. AS was supported a JP Morgan chase gift award and a UC Berkeley BAIR Commons Microsoft award.

REFERENCES

- [1] Rudolf Ahlswede and Andreas Winter. 2002. Strong converse for identification via quantum channels. *IEEE Trans. Inf. Theory* 48, 3 (March 2002), 569–579. <https://doi.org/10.1109/18.985947>
- [2] Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, and Umesh Vazirani. 2002. Dense quantum coding and quantum finite automata. *J. ACM* 49, 4 (July 2002), 496–511. <https://doi.org/10.1145/581771.581773>
- [3] Wojciech Banaszczyk. 1998. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Structures and Algorithms* 12, 4 (1998), 351–360. [https://doi.org/10.1002/\(SICI\)1098-2418\(199807\)12:4<351::AID-RSA3>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1098-2418(199807)12:4<351::AID-RSA3>3.0.CO;2-S)
- [4] Nikhil Bansal. 2010. Constructive Algorithms for Discrepancy Minimization. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, IEEE, 3–10. <https://doi.org/10.1109/focs.2010.7>
- [5] Nikhil Bansal. 2010. Constructive Algorithms for Discrepancy Minimization. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 3–10. <https://doi.org/10.1109/focs.2010.7>
- [6] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. 2019. An Algorithm for Komlós Conjecture Matching Banaszczyk's Bound. *SIAM J. Comput.* 48, 2 (Jan. 2019), 534–553. <https://doi.org/10.1137/17m1126795>
- [7] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. 2018. The Gram-Schmidt walk: A cure for the Banaszczyk blues. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 587–597. <https://doi.org/10.1145/3188745.3188850>
- [8] Imre Bárány. 2008. On the Power of Linear Dependencies. In *Bolyai Society Mathematical Studies*. Springer Berlin Heidelberg, 31–45. https://doi.org/10.1007/978-3-540-85221-6_1
- [9] Alexander Barvinok. 1995. Problems of distance geometry and convex properties of quadratic maps. *Discrete Comput. Geom.* 13, 2 (March 1995), 189–202. <https://doi.org/10.1007/bf02574037>
- [10] Joshua Batson, Daniel Spielman, and Nikhil Srivastava. 2012. Twice-Ramanujan Sparsifiers. *SIAM J. Comput.* 41, 6 (Jan. 2012), 1704–1721. <https://doi.org/10.1137/090772873>
- [11] Bernard Chazelle. 2000. *The Discrepancy Method*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511626371>
- [12] Daniel Dadush, Shashwat Garg, Shachar Lovett, and Aleksandar Nikolov. 2019. Towards a Constructive Version of Banaszczyk's Vector Balancing Theorem. *Theory of Computing* 15, 1 (2019), 1–58. <https://doi.org/10.4086/toc.2019.v015a015>
- [13] Daniel Dadush, Haotian Jiang, and Victor Reis. 2022. A New Framework for Matrix Discrepancy: Partial Coloring Bounds via Mirror Descent. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. ACM.
- [14] Ronen Eldan and Mohit Singh. 2018. Efficient algorithms for discrepancy minimization in convex sets. *Random Struct. Algor.* 53, 2 (Jan. 2018), 289–307. <https://doi.org/10.1002/rsa.20763>
- [15] Shashwat Garg. 2018. *Algorithms for combinatorial discrepancy*. Ph. D. Dissertation. Technische Universiteit Eindhoven.
- [16] Apostolos Giannopoulos. 1997. On some vector balancing problems. *Stud. Math.* 122, 3 (1997), 225–234. <https://doi.org/10.4064/sm-122-3-225-234>
- [17] Efim Gluskin. 1989. Extremal Properties Of Orthogonal Parallelepipeds And Their Applications To The Geometry Of Banach Spaces. *Math. USSR Sb.* 64, 1 (Feb. 1989), 85–96. <https://doi.org/10.1070/sm1989v06n01abeh003295>
- [18] David Gross. 2011. Recovering Low-Rank Matrices From Few Coefficients in Any Basis. *IEEE Trans. Inf. Theory* 57, 3 (March 2011), 1548–1566. <https://doi.org/10.1109/tit.2011.2104999>
- [19] Aram Harrow. 2013. The church of the symmetric subspace. *arXiv preprint arXiv:1308.6595* (2013).
- [20] Ilan Kremer, Noam Nisan, and Dana Ron. 1999. On Randomized One-round Communication Complexity. *Comput. Complex.* 8, 1 (June 1999), 21–49. <https://doi.org/10.1007/s00370050018>
- [21] Rasmus Kyng, Kyle Luh, and Zhao Song. 2020. Four deviations suffice for rank 1 matrices. *Adv. Math.* 375 (Dec. 2020), 107366. <https://doi.org/10.1016/j.aim.2020.107366>
- [22] Shachar Lovett and Raghu Meka. 2015. Constructive Discrepancy Minimization by Walking on the Edges. *SIAM J. Comput.* 44, 5 (Jan. 2015), 1573–1582. <https://doi.org/10.1137/130929400>
- [23] Adam Marcus, Daniel Spielman, and Nikhil Srivastava. 2015. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *Ann. Math.* (July 2015), 327–350. <https://doi.org/10.4007/annals.2015.182.1.8>
- [24] Jiri Matousek. 2009. *Geometric discrepancy: An illustrated guide*. Vol. 18. Springer Science and Business Media.
- [25] Raghu Meka. 2014. Discrepancy and beating the union bound. <https://windowsontimeory.org/2014/02/07/discrepancy-and-beating-the-union-bound/>
- [26] Ashwin Nayak. 1999. Optimal lower bounds for quantum automata and random access codes. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*. IEEE, IEEE Comput. Soc, 369–376. <https://doi.org/10.1109/sfcs.1999.814608>
- [27] Gábor Pataki. 1998. On the Rank of Extreme Matrices in Semidefinite Programs and the Multiplicity of Optimal Eigenvalues. *Math. Oper. Res.* 23, 2 (May 1998), 339–358. <https://doi.org/10.1287/moor.23.2.339>
- [28] Anup Rao and Amir Yehudayoff. 2020. *Communication Complexity*. Cambridge University Press. <https://doi.org/10.1017/cbo9781108671644>
- [29] Victor Reis and Thomas Rothvoss. 2020. Linear size sparsifier and the geometry of the operator norm ball. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2337–2348.
- [30] Thomas Rothvoss. 2014. Constructive Discrepancy Minimization for Convex Sets. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. Springer, IEEE, 380–391. <https://doi.org/10.1109/focs.2014.23>
- [31] Thomas Rothvoss. 2017. Constructive Discrepancy Minimization for Convex Sets. *SIAM J. Comput.* 46, 1 (Jan. 2017), 224–234. <https://doi.org/10.1137/141000282>
- [32] Joel Spencer. 1985. Six standard deviations suffice. *Trans. Amer. Math. Soc.* 289, 2 (1985), 679–706. <https://doi.org/10.1090/s0002-9947-1985-0784009-0>
- [33] Daniel Spielman and Nikhil Srivastava. 2011. Graph Sparsification by Effective Resistances. *SIAM J. Comput.* 40, 6 (Jan. 2011), 1913–1926. <https://doi.org/10.1137/080734029>
- [34] Armin Tavakoli, Alley Hameedi, Breno Marques, and Mohamed Bourennane. 2015. Quantum Random Access Codes Using Single d-Level Systems. *Phys. Rev. Lett.* 114, 17 (April 2015), 170502. <https://doi.org/10.1103/physrevlett.114.170502>
- [35] Anastasios Zouzias. 2012. A matrix hyperbolic cosine algorithm and applications. In *International Colloquium on Automata, Languages, and Programming*. Springer, 846–858.