

Spherical Discrepancy Minimization and Algorithmic Lower Bounds for Covering the Sphere

Chris Jones Matt McPartlon

June 14, 2020

Abstract

Inspired by the boolean discrepancy problem, we study the following optimization problem which we term SPHERICAL DISCREPANCY: given m unit vectors v_1, \dots, v_m , find another unit vector x that minimizes $\max_i \langle x, v_i \rangle$. We show that SPHERICAL DISCREPANCY is APX-hard and develop a multiplicative weights-based algorithm that achieves optimal worst-case error bounds up to lower order terms. We use our algorithm to give the first non-trivial lower bounds for the problem of covering a hypersphere by hyperspherical caps of uniform volume at least $2^{-o(\sqrt{n})}$. We accomplish this by proving a related covering bound in Gaussian space and showing that in this *large cap regime* the bound transfers to spherical space. Up to a log factor, our lower bounds match known upper bounds in the large cap regime.

1 Introduction

Let $S^{n-1} = \{x \in \mathbb{R}^n : \|x\|_2 = 1\}$ denote the surface of the sphere in \mathbb{R}^n . Suppose we have a collection of unit vectors $v_1, v_2, \dots, v_m \in S^{n-1}$. The goal of this work is to study the following optimization problem on the sphere, which we call SPHERICAL DISCREPANCY,

$$\min_{x \in S^{n-1}} \max_i \langle v_i, x \rangle.$$

The name comes from the boolean discrepancy problem in which x is required to be in $\{-1, +1\}^n$. The unit-norm requirement on x is crucial, otherwise the minimum is always either zero (achieved by the zero vector) or unbounded. SPHERICAL DISCREPANCY is a relaxation of the boolean discrepancy problem and a primary task of this paper is to adapt and improve upon algorithms from the boolean domain.

The SPHERICAL DISCREPANCY problem is intimately connected to the following covering problem on S^{n-1} : given m , what is the smallest value θ such that m spherical caps of angular radius θ can cover S^{n-1} ? Given a unit vector $v \in S^{n-1}$, corresponding to a pole, the cap associated with this pole is given by $\{x \in S^{n-1} : \langle v, x \rangle \geq \cos \theta\}$. Thus, a set of caps of angular radius θ covers the sphere if and only if the value of the SPHERICAL DISCREPANCY

instance on the poles is at least $\cos \theta$. This connection allows for a natural translation from algorithms for SPHERICAL DISCREPANCY to algorithmic lower bounds for the cap covering problem: given a sparse set of caps, by running the SPHERICAL DISCREPANCY algorithm on the poles we can produce a witness that lies outside of all the caps.

In this paper we develop an algorithm for SPHERICAL DISCREPANCY and use the above connection to prove the first non-trivial lower bounds for this sphere covering problem in what we call the *large cap regime*. In this regime, the volume of each spherical cap is required to be significantly large relative to the volume of the sphere. More precisely, each cap must cover a $2^{-\sqrt{n}}$ fraction of S^{n-1} . Outside of the large cap regime (and the regime in which the caps are very tiny), to the best of the authors' knowledge, no nontrivial lower bounds are known.

1.1 Prior Work on SPHERICAL DISCREPANCY

There are a few immediate algorithmic observations we can make about the SPHERICAL DISCREPANCY problem. If S^{n-1} is replaced by a convex body, SPHERICAL DISCREPANCY can be efficiently solved via standard convex programming techniques. In the case of the non-convex sphere, there are two simple exact algorithms with worst-case runtimes $m^{\Omega(n)}$: set up SPHERICAL DISCREPANCY directly as a quadratic program, or compute a spherical Voronoi diagram then search for a cell of maximal radius. Studying the problem under a different name, Petković et al [PPL12] develop a sophisticated recursive algorithm which is efficient in practice though it still has worst-case runtime $m^{\Omega(n)}$. There are additional algorithms and applications in \mathbb{R}^3 , see for example Cazals and Lorient [CL09]. For a number of applications with large n , see [PPL12].

It is NP-hard to construct a PTAS for SPHERICAL DISCREPANCY i.e. to output a solution whose value is within a factor $1 + \varepsilon$ of the true optimum (cf. Section 5). As in the boolean discrepancy problem, we are interested in $\text{poly}(m, n)$ -time approximation algorithms which achieve a worst-case bound independent of the true optimum value of the instance. With this goal in mind, it is natural to ask how well a uniformly random unit vector x performs. In the spherical setting, it is easy to show via a union bound that with high probability a vector $x \in_{\mathbb{R}} S^{n-1}$ achieves

$$\langle v_i, x \rangle \leq O\left(\sqrt{\frac{\ln m}{n}}\right)$$

for all i . We now compare this to the boolean version of the problem, and describe the improvements in that domain.

In the most general boolean discrepancy problem, we are given an m -by- n matrix A , and the goal is to minimize $\|Ax\|_{\infty}$ for $x \in \{\pm 1\}^n$. Representing the rows of the matrix by v_i , this is equivalent to minimizing the largest inner product $|\langle x, v_i \rangle|$. A particularly well-studied case of the boolean discrepancy problem assumes that A is the incidence matrix of a set system i.e. the entries are either from $\{-1, +1\}$ or $\{0, 1\}$. In this case picking a uniform $x \in_{\mathbb{R}} \{-1, +1\}^n$ is enough to achieve $|\langle v_i, x \rangle| \leq O(\sqrt{n \log m})$ for all i . A fundamental

result of discrepancy theory is Spencer’s “six standard deviations suffice” theorem, which says that this can be asymptotically improved to $O(\sqrt{n \log \frac{m}{n}})$,

Theorem 1 [Spe85]. *Given $v_1, \dots, v_m \in \{0, 1\}^n$ with $m \geq n$, there is $x \in \{-1, +1\}^n$ such that*

$$|\langle v_i, x \rangle| \leq 6\sqrt{n \log \left(\frac{m}{n} + 1 \right)}$$

for all i .

A famous open problem in discrepancy theory, the Komlós conjecture, states that (in the case $m = n$) a similar bound holds if we relax the set system assumption to an assumption that the column norms of A are at most \sqrt{n} (or after normalization, at most 1),

Conjecture 1. *Let $w_1, \dots, w_n \in \mathbb{R}^n$ be vectors with $\|w_i\|_2 \leq 1$, and arrange the w_i as columns of a matrix W . There is a boolean vector $x \in \{+1, -1\}^n$ so that $\|Wx\|_\infty = O(1)$.*

When the row norms of A are also assumed to have norm at most 1 (which is often without much loss of generality, see Section 4), this is the setting in which SPHERICAL DISCREPANCY is a relaxation of the boolean discrepancy problem. Furthermore, we will work with the “one-sided” discrepancy $\langle v_i, x \rangle$ rather than using the absolute value of the inner product. This difference is mostly inconsequential as one can throw in $-v_i$ to bound the absolute value. The absence of absolute value corresponds more naturally to the geometric questions we consider in Section 3.

Spencer’s argument is evidently an improvement on the random bound in the case $m = O(n)$. However, if $m \geq n^{1+c}$ for some $c > 0$, then the qualitative bound $O(\sqrt{n \log \frac{m}{n}})$ is equivalent to $O(\sqrt{n \log m})$ – the same as in the random case. Thus one must look into the actual constants and error terms to see if Spencer’s bound beats the random bound whenever m is significantly superlinear in n . In fact, the precise constant in the random argument is $\sqrt{2n \ln m} \cdot (1 + o(1))$, which is definitely stronger than $6\sqrt{n \log \frac{m}{n}}$ when m is superpolynomial in n ; we are not sure the exact regime of m in which Spencer’s analysis (or any of the follow-up reproofs of the result) beats the random argument. These issues of constant factors will plague us in the spherical case, where we must prove tight results all the way up to $m = 2^{\sqrt{n}}$.

Spencer’s theorem was only recently made algorithmic, and in the last few years there has been a spate of recent activity on algorithmic solutions to problems in discrepancy theory [Ban12, LM15, BDG19, BDGL18, LRR17]. Our algorithm for SPHERICAL DISCREPANCY is based on a deterministic algorithm for Spencer’s theorem due to Levy, Ramadas, and Rothvoss [LRR17], which itself is a derandomization of a random walk-based algorithm of Lovett and Meka [LM15]. The Lovett-Meka algorithm is part of a general class of *partial coloring* algorithms for boolean discrepancy which produce a coloring $x \in \{-1, +1\}^n$ in $\log n$ rounds. In each round, half the remaining coordinates are set to $+1$ or -1 . Our basic strategy for the SPHERICAL DISCREPANCY problem is to take the first round only of a partial coloring algorithm, as this produces a vector with large norm; there is no need to further round it towards a corner of the boolean hypercube.

Finally, we mention vector discrepancy, a different relaxation of boolean discrepancy used by Lovasz [Lov00]. Fix an m -by- n matrix A . Instead of assigning $\{\pm 1\}$ to entries of a vector x , we assign unit vectors x_i in some larger dimension, with the goal to minimize

$$\max_{i=1,\dots,m} \left\| \sum_{j=1}^n A_{ij} x_i \right\|_2.$$

When the dimension of x_i is n , the problem is convex, and recently Nikolov [Nik13] was able to verify the Komlós conjecture in this setting using techniques from convex programming. In the spherical setting we have a global constraint $\sum_{i=1}^n x_i^2 = 1$ whereas vector discrepancy (like most other convex relaxation techniques) relaxes each variable independently. Practically speaking, the critical difference between the spherical and vector relaxations is that the domain of spherical discrepancy is non-convex and so we no longer have access to powerful tools such as duality.

For more background about discrepancy theory, see the book by Chazelle [Cha00].

1.2 Prior Work on Sphere Covering Lower Bounds.

The *spherical cap* with pole $v \in S^{n-1}$ and angular radius θ is the set

$$\{x \in S^{n-1} \mid \langle x, v \rangle \geq \cos \theta\}.$$

The normalized volume of a measurable set $C \subseteq S^{n-1}$ is

$$\text{vol}(C) \stackrel{\text{def}}{=} \Pr_{x \in_{\mathbb{R}} S^{n-1}} [x \in C].$$

There are two dual questions to ask about spherical caps. The *packing* question asks: given m , what is the largest δ so that m spherical caps of normalized volume δ can be arranged disjointly in S^{n-1} ? The *covering* question asks: given m , what is the smallest δ so that there are m spherical caps of normalized volume δ which cover S^{n-1} ? For both of these questions, a trivial *volume bound* applies: if m caps of normalized volume δ cover (respectively pack) S^{n-1} , then necessarily $\delta \geq 1/m$ (respectively $\delta \leq 1/m$). This could only be achieved if the caps could be arranged disjointly on the surface of the sphere, which is impossible except for the case of two hemispheres. The quantity δm is called the density of the packing/covering.

The study of these questions originates from the study of maximum density packings/minimum density coverings of spheres in \mathbb{R}^n . The pioneering work of Rogers and coauthors [Rog58, CFR59] led to the *simplex bound*, which states that the density of a packing or covering cannot beat a natural strategy based on tiling \mathbb{R}^n with a regular simplex (note that \mathbb{R}^n for $n \geq 3$ cannot be tiled with a regular simplex, so this bound is not tight).

In spherical space (and hyperbolic space), the simplex bound for packing was extended by Böröczky [Bör78]. For covering in spheres, we have very few bounds outside of the trivial volume bound, and the simplex bound for covering, stated implicitly in [CFR59] and explicitly in [Bör04, Conjecture 6.7.3], has remained unproven,

Conjecture 2. *If a set of m spherical caps each of normalized volume $\delta < 1/2$ covers S^{n-1} , then*

$$m\delta \geq \tau_{n,\delta}$$

where $\tau_{n,\delta}$ is defined for the interested reader in Section 6. For δ sufficiently small, the conjectured density is $\tau_{n,\delta} \approx \frac{n}{e\sqrt{e}}$ i.e. a factor of $\Omega(n)$ higher than the trivial volume bound.¹ It should be noted that in general there is no perfect relationship between packing and covering, for example taking an optimal packing and extending the caps just enough to cover the whole of S^{n-1} will not result in an optimal covering; there are coverings with smaller density [SV05]. In general much more is known about packings than coverings across Euclidean, spherical, and hyperbolic spaces.

Conjecture 2 has been shown to hold for $n = 3$ [Bör04, Theorem 5.1.1]. Conjecture 2 is tight when S^{n-1} can be tiled with regular spherical simplices using m vertices, corresponding to the projection to S^{n-1} of an n -dimensional $\{3, 3, \dots, p\}$ Coxeter polytope such as the regular simplex ($m = n + 1$) and the cross-polytope ($m = 2n$), as observed by Coxeter [Cox63].

Conjecture 2 has also been confirmed in the regime where δ is small enough that the caps have angular radius $\theta \leq \frac{1}{\sqrt{n}}$. In this regime, the caps are small enough that S^n looks very similar to \mathbb{R}^n , and the lower bound techniques used by Coxeter-Rogers-Few [CFR59] in \mathbb{R}^n are enough to verify Conjecture 2. The authors of [CFR59] note that their technique can be extended to S^n , but they do not analyze the full range of δ for which their proof goes through. Their result only holds when $\theta \leq 1/\sqrt{n}$; see [Bör04, Lemma 6.8.4] for an exposition of the proof.

There is also a nontrivial covering lower bound in the regime where $\delta \geq 1/n$, which corresponds to caps of angular radius $\frac{\pi}{2} - \Theta(n^{-1/2})$. The lower bound comes from the Lusternik-Schnirelmann theorem,

Theorem 2. *If S^{n-1} is covered by n open or closed sets, then one of those sets contains a pair of antipodal points.*

Since a spherical cap with any volume $\delta < 1/2$ does not contain antipodal points, any cover must use at least $m \geq n + 1$ caps. When $\delta = \Omega(1)$ this translates to a linear lower bound on the density $m\delta = \Omega(n)$, which matches Conjecture 2 up to a constant. In the 60 years since [CFR59], we are aware of no covering lower bounds better than the trivial $m\delta \geq 1$ outside of these two extremes.

On the flip side, there are constructions of spherical cap covers that nearly match Conjecture 2,

Theorem 3 [BW03]. *For any $0 < \varphi \leq \arccos \frac{1}{\sqrt{n}}$, there is an arrangement of spherical balls with radius φ with density at most*

$$c \cdot n \ln(1 + (n - 1) \cos^2 \varphi)$$

¹We were unable to find a reference for the seemingly obvious fact that $\tau_{n,\delta} \geq c \cdot n$ for some positive constant c . See Conjecture 5.

where c is an absolute constant.

Translating the assumption on φ to the parameterization by δ , the upper bound on φ says that δ is less than some fixed constant. The density bound achieved by the theorem is $O(n)$ for $\delta = \Theta(1)$ and increases to $O(n \log n)$ for constant θ . The authors provide another quasilinear bound with a tighter constant,

Theorem 4 [BW03]. *For any $\theta < \frac{\pi}{2}$, there is a covering by spherical caps of radius θ with density at most*

$$n \ln n + c \cdot n \ln \ln n$$

where c is an absolute constant.

We also point out the relationship with hitting sets for spherical caps. A set of points $P \subset S^{n-1}$ is a *hitting set* for spherical caps of volume δ if for every spherical cap C with volume δ , we have $P \cap C \neq \emptyset$. Observe that a hitting set is the same as a cover by spherical caps. The constructions above yield hitting sets for caps of volume δ of size $\tilde{O}(n/\delta)$ — however, they are randomized. Rabani and Shpilka [RS09] show how to deterministically construct a hitting set of size $\text{poly}(n, 1/\delta)$. Their construction works only in the large cap regime, $\delta \geq 2^{-\Omega(\sqrt{n})}$.

1.3 Our Results and Organization of the Paper.

In Section 2 we give an algorithm for SPHERICAL DISCREPANCY which, analogously to Spencer’s theorem, improves upon the worst-case guarantee of the random algorithm,

Theorem 5. *Let $m \geq 16n$ and $v_1, v_2, \dots, v_m \in S^{n-1}$ be unit vectors. We can find a vector $x \in S^{n-1}$ such that*

$$\langle v_i, x \rangle \leq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 + O\left(\frac{1}{\log \frac{m}{n}}\right) \right)$$

for all i , via a deterministic algorithm that runs in time polynomial in m and n .

The algorithm behind Theorem 5 is based on a recent deterministic multiplicative weights algorithm for boolean discrepancy due to Levy, Ramadas, and Rothvoss [LRR17]. The crux of the above theorem is that we are able to achieve the optimal constant (but not optimal error term), for a large range of m up to essentially $2^{\sqrt{n}}$,

Theorem 6. *For every choice of $2^{-o(\sqrt{n})} < \delta < \frac{1}{n^2}$, there is a parameter $1/\delta \leq m \leq \frac{2n \ln n}{\delta}$ and unit vectors $v_1, \dots, v_m \in S^{n-1}$ such that, for any $x \in S^{n-1}$ there is a v_i with*

$$\langle x, v_i \rangle \geq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 - o\left(\frac{1}{\sqrt{\log \frac{m}{n}}}\right) \right).$$

The proof of Theorem 6 is in Section 3. In that section, we work out the consequences of the relationship between the SPHERICAL DISCREPANCY problem and covering the sphere by

spherical caps. Our primary result is an application of Theorem 5 to prove the simplex bound for spherical space (Conjecture 2) up to a log factor for caps that are not too small,²

Theorem 7. *If a set of m spherical caps each of normalized volume $2^{-o(\sqrt{n})} \leq \delta < \frac{1}{2}$ covers S^{n-1} , then*

$$\delta \geq \Omega \left(\frac{n}{m \sqrt{\log(1 + \frac{m}{n})}} \right).$$

We are able to beat the trivial volume lower bound $m\delta \geq 1$ in this theorem exactly when the algorithm behind Theorem 5 beats the trivial random bound. This requires us to carefully optimize the performance of the algorithm. The lower bound is algorithmic in the sense that, given a list of the m caps of volume δ that does not meet the bound, there is a $\text{poly}(n, m)$ time algorithm that finds a point outside all caps. The lower bound assumes that the spherical caps have normalized volume at least $2^{-\sqrt{n}}$, a setting we term the *large cap regime*. We explain in Section 3 how the assumption $\delta \geq 2^{-\sqrt{n}}$ naturally corresponds to a geometric property of this regime of δ , namely that caps of this volume have approximately the same spherical measure as the Gaussian measure of their corresponding halfspaces. When this approximation holds, the lower bound follows as a corollary of the following covering lower bound for Gaussian space, which itself follows naturally from Theorem 5,

Theorem 8. *If m halfspaces of Gaussian measure $\delta < \frac{1}{2}$ cover the \sqrt{n} -sphere in \mathbb{R}^n , then*

$$\delta \geq \Omega \left(\frac{n}{m \sqrt{\log(1 + \frac{m}{n})}} \right).$$

The details of the proof are found in Section 3. Finally, in Section 3 we also study the performance of the following algorithm for generating sphere packings: run the algorithm behind Theorem 5 for m iterations. We show

Theorem 9. *Let $16n \leq m \leq 2^{o(\sqrt{n})}$. There is a deterministic algorithm which runs in time $\text{poly}(m, n)$ that outputs m points $v_1, v_2, \dots, v_m \in S^{n-1}$ such that*

$$\langle v_i, v_j \rangle \leq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 + O \left(\frac{1}{\log \frac{m}{n}} \right) \right)$$

for all pairs i, j . Taking the maximal radius r such that spherical caps of radius r around the v_i are disjoint produces a packing with density $\Omega \left(\frac{n}{2^n \sqrt{\log \frac{m}{n}}} \right)$.

In Section 4 we show that the natural relaxation of the Komlós conjecture for spherical discrepancy is true,

Theorem 10. *Let $w_1, \dots, w_n \in \mathbb{R}^n$ be vectors with $\|w_i\|_2 \leq 1$, and let W be the matrix with the w_i as columns. Then we can find a unit vector $x \in S^{n-1}$ such that*

$$\|Wx\|_\infty = O \left(\frac{1}{\sqrt{n}} \right)$$

²The “log factor” becomes $n^{1/4}$ when m is near $2^{\sqrt{n}}$.

in time polynomial in n .

The algorithm in Section 2 can be modified to show this, and we explain how to do this, though for the proof as pointed out by the anonymous referee it suffices to simply cite the partial coloring lemma from boolean discrepancy.

In Section 5 we show that the SPHERICAL DISCREPANCY problem is APX-hard via a gap-preserving reduction from NAE-3-SAT,

Theorem 11. *There is a constant $C > 1$ so that it is NP-hard to distinguish between instances of SPHERICAL DISCREPANCY with $m = O(n)$ with value at most $\frac{1}{\sqrt{n}}$, and instances with value at least $\frac{C}{\sqrt{n}}$.*

Lastly, in Section 6 we conclude with some further questions.

All asymptotic notation in this paper is in terms of n , while different parameters such as m and δ are functions of n . The functions \exp and \ln denote base e , whereas \log is used to denote base 2.

2 Multiplicative Weights for SPHERICAL DISCREPANCY.

In this section we develop the algorithm for Theorem 5,

Theorem 5. *Let $m \geq 16n$ and $v_1, v_2, \dots, v_m \in S^{n-1}$ be unit vectors. We can find a vector $x \in S^{n-1}$ such that*

$$\langle v_i, x \rangle \leq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 + O\left(\frac{1}{\log \frac{m}{n}}\right) \right)$$

for all i , via a deterministic algorithm that runs in time polynomial in m and n .

It is critical in our applications to have both the small error term and the constant $\sqrt{2}$ in the bound of Theorem 5. In the setting of boolean discrepancy, the partial coloring method introduces a flexible constant and so it is difficult in that setting to achieve the correct constant. In fact, as pointed out by the anonymous reviewer, an application of the partial coloring lemma can be used to prove Theorem 5 with the weaker bound $O\left(\sqrt{\frac{\log \frac{m}{n}}{n}}\right)$.

We do not give the details though it essentially matches the argument in Section 4. In the spherical setting, we are able to skirt partial coloring because the set of “colorings” in our domain is S^{n-1} , which can be rounded to by simply normalizing a “candidate coloring”.

The algorithm behind Theorem 5 is based on a recent deterministic multiplicative weights-based algorithm for boolean discrepancy due to Levy, Ramadas, and Rothvoss [LRR17]. In the general setting of the multiplicative weight update (MWU) method, we are trying to make repeated decisions using the opinions of a collection of “experts”; some of the experts tell us good advice whereas others may not, and we must learn a combination of the experts which collectively leads to a wise decision. Though the analogy is not perfect, in our case the decision is a small update direction for a candidate coloring x . Multiplicative weights

can often be used to obtain an upper bound on the worst-case number of “mistakes” from following any of the experts, and in our case that corresponds to a bound on the worst-case inner product with any of the v_i . For more on multiplicative weights, see the survey paper [AHK12].

Similarly to their algorithm, we slowly grow a candidate coloring x which is initially zero, maintaining a multiplicative weight for each input vector v_i to help choose the next update direction. The weights w_i essentially equal $\exp(\langle x, v_i \rangle)$ (Lemma 2.4). We choose an update direction for x from an eigenspace corresponding to a small eigenvalue of the matrix $\sum_{i=1}^m w_i v_i v_i^\top$, subject to lying in a particular subspace of (linear) constraints. Ideally, we would like to move in the smallest eigenspace of this matrix. For example, consider the case when the vectors v_i are mutually orthogonal. If at each timestep we make a small update in the direction of the v_i with the current smallest weight, and then update the weights, this will ensure all weights remain approximately equal. Stopping the algorithm when $\|x\|_2 = 1$ will produce x with the optimal $\frac{1}{\sqrt{n}}$ inner product with each vector. However, in general it is impossible to move in the smallest eigenspace and also the subspace of constraints. Improving on [LRR17], we are able to obey the constraints yet move in a direction that is essentially no worse than the average eigenvalue (Lemma 2.2). Under the heuristic assumption that the weights remain somewhat balanced during the run of the algorithm, this is nearly as good as moving in the smallest eigenspace.³

The proof of correctness also follows the lines of Levy, Ramadas, and Rothvoss, though few technical details in the algorithm and the proof remain the same. We must do a careful induction to show that our algorithm is well-defined, and a careful analysis and choice of parameters to ensure it meets Theorem 5.

2.1 The Algorithm.

Let v_1, \dots, v_m be unit vectors in \mathbb{R}^n . Note that the guarantee of the theorem is trivial for $m \geq n \exp(n/2)$. If the input m is this large, we return any unit vector. For the remainder of this algorithm we assume that $m \leq n \exp(n/2)$.

For a PSD matrix $M \in \mathbb{R}^{n \times n}$, denote the eigenvalues of M by $\mu_1(M) \geq \mu_2(M) \geq \dots \geq \mu_n(M)$ with corresponding eigenvectors $u_i(M)$ to $\mu_i(M)$. We denote by S^\perp the orthogonal complement of the span of a set of vectors S .

We will use parameters λ, ρ, δ , and T .

Description of the parameters:

- $\lambda = \sqrt{\ln \frac{m}{n}}$
- $\delta = \frac{1}{n^3}$ is the step size for updates to $x^{(t)}$
- $T = \frac{2(n-5)}{\delta^2}$ is the number of iterations

³The weights won't remain somewhat balanced, because for example two v_i may point in opposite directions. But we hope that the “critical” weights will be somewhat balanced.

MWU for Spherical Discrepancy

Input: unit vectors $v_1, \dots, v_m \in S^{n-1}$

$x^{(0)} \leftarrow 0^n$

$w_i^{(0)} \leftarrow \exp(-\lambda^2)$

for $t = 0, \dots, T$:

$I^{(t)} \leftarrow \{i \in [m] : w_i^{(t)} \geq 2\}$

$M^{(t)} \leftarrow \sum_{i \notin I} w_i^{(t)} v_i v_i^\top$

$P^{(t)} \leftarrow \{x^{(t)}\} \cup \{\sum_{i \in [m]} w_i^{(t)} v_i\} \cup \{v_i : i \in I^{(t)}\} \cup \{u_j(M^{(t)}) : 1 \leq j \leq n - |I^{(t)}| - 3\}$

$y^{(t)} \leftarrow$ any unit vector in $P^{(t)\perp}$

$x^{(t+1)} \leftarrow x^{(t)} + \delta y^{(t)}$

for $i = 1, \dots, m$:

$w_i^{(t+1)} \leftarrow w_i^{(t)} \cdot \exp(\lambda \langle v_i, \delta y^{(t)} \rangle) \cdot \rho$

return $x^{(T)} / \|x^{(T)}\|_2$

- $\rho = \exp\left(-\frac{\delta^2 \lambda^2}{2(n-5)} \cdot (1 + \lambda \delta n)\right)$ is the discount factor in the weight update step. ρ is very slightly less than 1.

The initial weights $w_i^{(0)}$ don't affect the output of the algorithm so long as they're uniform.

2.2 Runtime Analysis

The vectors v_i should be specified to $2 \log n$ bits of precision so that the error in $\langle v_i, x \rangle$ can be incorporated into the error term of Theorem 5.

In addition, all numerical calculations should be truncated to $30 \log n$ bits of precision. The statements of all lemmas can be modified to include a polynomially small error term, which ultimately does not affect the statement of Theorem 5. As an example, though the value ρ is close to 1, it has only $\Theta(\log n)$ zeros after the decimal point, and truncating to $30 \log n$ bits is enough to approximate it throughout the $T = O(n^7)$ iterations of the algorithm.

Each iteration takes $O(mn^2 + n^3)$ time: evaluating M takes time mn^2 , and computing an eigendecomposition of $M^{(t)}$ can be done in time $O(n^3)$. There are $T = O(n^7)$ iterations, for an overall runtime of $O(mn^9 + n^{10})$.

It seems likely that fewer iterations T are sufficient (we must set the corresponding $\delta = \sqrt{\frac{2(n-5)}{T}}$). Matrix multiplication methods could also be used to lower the exponent in the time needed to compute an eigendecomposition, but we don't optimize the runtime here.

2.3 Bounding the Maximum Inner Product

The analysis of the algorithm will use the potential function $\Phi(t)$,

$$\Phi(t) \stackrel{\text{def}}{=} \sum_{i=1}^m w_i^{(t)}.$$

Initially, the potential function is $\Phi(0) = \sum_{i=1}^m \exp(-\lambda^2) = \sum_{i=1}^m \frac{n}{m} = n$.

We choose the discount factor ρ so that $\Phi(t+1) \leq \Phi(t)$ for every t . That is, after setting $x^{(t+1)} \leftarrow x^{(t)} + \delta y^{(t)}$, each weight is increased by a factor proportional to $\exp(\lambda \cdot \delta \langle v_i, y^{(t)} \rangle)$, which seems like it could increase the potential; ρ is chosen just small enough to counteract the increase. The key to producing a tight bound on $\langle v_i, x^{(T)} \rangle$ lies in maximizing ρ while still ensuring that the potential is decreasing.

As written, it is not clear that the algorithm is well-defined; it is a priori possible that the space $P^{(t)\top}$ is trivial and does not contain a unit vector. We say that the algorithm *succeeds up to time t* if the following conditions occur:

- (i) $P^{(t_0)\perp} \neq 0$ for every $t_0 < t$ (and therefore all weights $w_i^{(t_0)}$ are properly defined as are $w_i^{(t)}$).
- (ii) The potential function Φ is nonincreasing from time 0 to time t .

We now prove that if the algorithm succeeds up to time t , it succeeds up to time $t+1$, and so by induction the algorithm succeeds up to time T .

Lemma 2.1. *If $|I^{(t)}| \leq n-3$, then $P^{(t)\perp} \neq 0$.*

Proof. We need to check that $\dim(\text{span}(P^{(t)})) < n$. Of the four sets composing $P^{(t)}$, the first three contain at most $|I^{(t)}| + 2$ vectors. Under the assumption on $|I^{(t)}|$ the fourth set contains at most $n - |I^{(t)}| - 3$ vectors and so the result follows. \square

If the potential function is nonincreasing up to time t , then $\Phi(t) \leq n$. Therefore, on iteration $t+1$ the set $I^{(t)}$ has size at most $\frac{n}{2}$, and by the lemma just above $P^{(t)\perp}$ contains a unit vector and the next set of weights will be well-defined. This fulfills the first condition for success up to time $t+1$. It remains to work towards the second condition. We at least know that $\Phi(t+1)$ is well-defined because the weights $w_i^{(t+1)}$ are well-defined — we just need to prove $\Phi(t+1) \leq \Phi(t)$.

Lemma 2.2. *For each unit vector $y \in P^{(t)\perp}$, one has $y^\top M^{(t)} y \leq \frac{\Phi(t)}{n-5}$*

Proof. Because of the fourth set composing $P^{(t)}$, y is perpendicular to many eigenvalues of $M^{(t)}$, so we bound the max eigenvalue of the remaining eigenspaces. Recall that

$$M^{(t)} = \sum_{i \notin I^{(t)}} w_i^{(t)} v_i v_i^\top.$$

For any $i \in [m]$, we have $\text{tr}(v_i v_i^\top) = \|v_i\|^2 = 1$, thus

$$\text{tr}(M^{(t)}) = \sum_{i \notin I^{(t)}} w_i^{(t)} = \sum_{i=1}^m w_i^{(t)} - \sum_{i \in I^{(t)}} w_i^{(t)} \leq \Phi(t) - 2|I^{(t)}|.$$

On the other hand $\text{tr}(M^{(t)})$ is the sum of the eigenvalues of $M^{(t)}$. By a use of Markov's inequality, the $(n - |I^{(t)}| - 2)^{\text{th}}$ largest eigenvalue of $M^{(t)}$ is at most $\frac{\Phi(t) - 2|I^{(t)}|}{n - |I^{(t)}| - 2}$. Symbolically,

$$\mu_{n-|I^{(t)}|-2} \leq \frac{\Phi(t) - 2|I^{(t)}|}{n - |I^{(t)}| - 2} = \frac{\Phi(t)}{n - 5} \cdot \frac{(n - 5) \left(1 - \frac{|I^{(t)}|}{n}\right)}{n - |I^{(t)}| - 2} \cdot \frac{1 - \frac{2|I^{(t)}|}{\Phi(t)}}{1 - \frac{|I^{(t)}|}{n}}$$

From the inductive assumption that $\Phi(t)$ is nonincreasing, $\Phi(t) \leq \Phi(0) = n$, and therefore the right term is bounded by 1. The middle term is

$$\frac{n - |I^{(t)}| - 5 + \frac{5|I^{(t)}|}{n}}{n - |I^{(t)}| - 2}$$

The bound $|I^{(t)}| \leq n/2$ shows that this is also at most 1. Hence $\mu_{n-|I^{(t)}|-2} \leq \frac{\Phi(t)}{n-5}$. Since y is a linear combination of eigenvectors with eigenvalues at most this value, we have $y^\top M^{(t)} y \leq \frac{\Phi(t)}{n-5}$. \square

Fact 2.1. For any $0 \leq x \leq 1$,

$$e^x \leq 1 + x + \frac{x^2}{2} + \frac{x^3}{2}.$$

Proof. See Appendix A. \square

Lemma 2.3. $\Phi(t+1) \leq \Phi(t)$.

Proof. The recursive update for the weights at time $t+1$ is given in the algorithm,

$$\Phi(t+1) = \sum_{i=1}^m w_i^{(t+1)} = \sum_{i=1}^m w_i^{(t)} \cdot \exp(\lambda \delta \langle v_i, y^{(t)} \rangle) \cdot \rho.$$

We assumed that $m \leq n \exp(n/2)$ so that $\lambda \leq \sqrt{n/2}$ and $\lambda \delta \langle v_i, y^{(t)} \rangle \leq \lambda \delta \leq 1$, and we can apply the previous Lemma,

$$\begin{aligned} \Phi(t+1)/\rho &\leq \sum_{i=1}^m w_i^{(t)} \left(1 + \lambda \delta \langle v_i, y^{(t)} \rangle + \frac{\lambda^2 \delta^2 \langle v_i, y^{(t)} \rangle^2}{2} + \frac{\lambda^3 \delta^3 \langle v_i, y^{(t)} \rangle^3}{2} \right) \\ &= \sum_{i=1}^m w_i^{(t)} + \lambda \delta \left\langle \sum_{i=1}^m w_i^{(t)} v_i, y^{(t)} \right\rangle + \frac{\lambda^2 \delta^2}{2} \sum w_i^{(t)} \langle v_i, y^{(t)} \rangle^2 + \frac{\lambda^3 \delta^3}{2} \sum w_i^{(t)} \langle v_i, y^{(t)} \rangle^3. \end{aligned}$$

Since $y^{(t)} \in P^{(t)\perp}$, the second term is zero. Furthermore, $\sum_{i=1}^m w_i^{(t)} \langle v_i, y^{(t)} \rangle^2 = y^{(t)\top} M^{(t)} y^{(t)}$, so by Lemma 2.2 we have,

$$\begin{aligned} \Phi(t+1)/\rho &\leq \Phi(t) + \frac{\lambda^2 \delta^2}{2(n-5)} \Phi(t) + \frac{\lambda^3 \delta^3}{2} \sum w_i^{(t)} \underbrace{\langle v_i, y^{(t)} \rangle^3}_{\leq 1} \\ &\leq \Phi(t) + \frac{\lambda^2 \delta^2}{2(n-5)} \Phi(t) + \frac{\lambda^3 \delta^3}{2} \Phi(t) \\ &= \Phi(t) \cdot \left(1 + \frac{\lambda^2 \delta^2}{2(n-5)} (1 + \lambda \delta n) \right) \\ \Phi(t+1)/\rho &\leq \Phi(t) \cdot \exp\left(\frac{\lambda^2 \delta^2}{2(n-5)} (1 + \lambda \delta n) \right). \end{aligned}$$

In the last line, we use the inequality $1 + x \leq e^x$. The exponential term is exactly $1/\rho$, therefore we conclude $\Phi(t+1) \leq \Phi(t)$. \square

This finishes the proof by induction that the algorithm succeeds up to time T . We now proceed to bound the max inner product $\langle v_i, x^{(T)} \rangle$, starting with a few lemmas.

Lemma 2.4. $w_i^{(t)} = \exp(\lambda \langle v_i, x^{(t)} \rangle - \lambda^2) \cdot \rho^t$

Proof. The weights are initially $\exp(-\lambda^2)$. Each iteration they are multiplied by a factor of ρ , and also by $\exp(\lambda \delta \langle v_i, y^{(t)} \rangle)$, so

$$w_i^{(t)} = \exp\left(\lambda \left\langle v_i, \sum_{t' < t} \delta y^{(t')} \right\rangle - \lambda^2 \right) \cdot \rho^t.$$

On the other hand, $x^{(t)} = \sum_{t' < t} \delta y^{(t')}$. \square

Lemma 2.5. $\|x^{(t)}\|_2 = \delta \sqrt{t}$

Proof. Note that $x^{(t)} = x^{(t-1)} + \delta y^{(t-1)}$, and since $y^{(t)}$ is a unit vector orthogonal to $x^{(t)}$ due to the subspace U_1 ,

$$\|x^{(t)}\|_2^2 = \|x^{(t-1)}\|_2^2 + \delta^2 \|y^{(t-1)}\|_2^2 = \|x^{(t-1)}\|_2^2 + \delta^2.$$

Since $x^{(0)} = \mathbf{0}^n$, the result follows. \square

Lemma 2.6. *At all times t , the weights satisfy $\max_i w_i^{(t)} \leq 3$.*

Proof. Once a weight becomes greater than 2, it moves into the set $I^{(t)}$ and further moves are orthogonal to v_i , meaning the weight only decreases due to ρ . On the update that the weight moves into $I^{(t)}$,

$$w_i^{(t)} = w_i^{(t-1)} \cdot \exp(\lambda \cdot \delta \langle v_i, y^{(t-1)} \rangle) \cdot \rho$$

$$\begin{aligned} &\leq 2 \cdot \exp(\lambda \cdot \delta) \cdot \rho \\ &\leq 2 \cdot \exp(\lambda \cdot \delta) \leq 3 \end{aligned}$$

Note that the weight may move out of and back into $I^{(t)}$ in the future due to the discount factor ρ . \square

We are now in a position to prove Theorem 5. We will take x to be $x^{(T)}/\|x^{(T)}\|$.

Lemma 2.7. $\langle v_i, x^{(T)}/\|x^{(T)}\| \rangle \leq \lambda\sqrt{\frac{2}{n}} + O\left(\frac{1}{\lambda\sqrt{n}}\right)$

Proof. We have an exact expression for the weights by Lemma 2.4 and a bound on the weights by Lemma 2.6,

$$w_i^{(T)} = \exp(\lambda \langle v_i, x^{(T)} \rangle - \lambda^2) \cdot \rho^T \leq 3.$$

We take logs, solve for $\langle v_i, x^{(T)} \rangle$, and optimize the resulting bound.

$$\begin{aligned} \lambda \langle v_i, x^{(T)} \rangle - \lambda^2 + T \ln \rho &\leq \ln 3 \\ \langle v_i, x^{(T)} \rangle &\leq \lambda - \frac{T \ln \rho}{\lambda} + \frac{1}{\lambda} \ln 3 \end{aligned}$$

Plug in $\ln \rho = -\frac{\lambda^2 \delta^2}{2(n-5)} \cdot (1 + \lambda \delta n)$ and normalize by $\|x^{(T)}\| = \delta\sqrt{T}$ (Lemma 2.5),

$$\begin{aligned} \langle v_i, x^{(T)}/\|x^{(T)}\| \rangle &\leq \frac{\lambda}{\delta\sqrt{T}} + \frac{\lambda\delta\sqrt{T}}{2(n-5)} \cdot (1 + \lambda\delta n) + \frac{1}{\lambda\delta\sqrt{T}} \ln 3 \\ &\leq \left(\frac{\lambda}{\delta\sqrt{T}} + \frac{\lambda\delta\sqrt{T}}{2(n-5)} \right) (1 + \lambda\delta n) + \frac{1}{\lambda\delta\sqrt{T}} \ln 3. \end{aligned}$$

The choice of parameters $\delta\sqrt{T} = \sqrt{2(n-5)}$ balances the first two terms to the value $\frac{\lambda}{\sqrt{2(n-5)}}$.

$$\langle v_i, x^{(T)}/\|x^{(T)}\| \rangle \leq \lambda\sqrt{\frac{2}{n-5}}(1 + \lambda\delta n) + O\left(\frac{1}{\lambda\sqrt{n}}\right)$$

Recalling that $\lambda \leq \sqrt{n}$ and $\delta = \frac{1}{n^3}$, we can absorb the term $\lambda\delta n$ into the remaining error,

$$= \lambda\sqrt{\frac{2}{n-5}} + O\left(\frac{1}{\lambda\sqrt{n}}\right).$$

Bounding $\sqrt{\frac{1}{n-5}} \leq \sqrt{\frac{1}{n}} + \frac{5}{n^{3/2}}$ shows we can incorporate this term into the error as well. \square

Finally, if we plug in $\lambda = \sqrt{\log \frac{m}{n}}$ to Lemma 2.7 we recover Theorem 5.

3 Lower Bounds for Covering Problems.

As described in the introduction, algorithms for SPHERICAL DISCREPANCY immediately yield algorithmic lower bounds for the problem of covering the hypersphere by hyperspherical caps, which will let us show the following theorem,

Theorem 7. *If a set of m spherical caps each of normalized volume $2^{-o(\sqrt{n})} \leq \delta < \frac{1}{2}$ covers S^{n-1} , then*

$$\delta \geq \Omega\left(\frac{n}{m\sqrt{\log(1 + \frac{m}{n})}}\right).$$

Essentially, the proof boils down to the following calculation: set the right-hand side of Theorem 5 equal to $\cos \theta$ and solve for the normalized volume δ of a spherical cap with angular radius θ . The calculation is simplified by noticing that Theorem 5 has a form which applies perfectly to the following question about Gaussian space: how many halfspaces of Gaussian measure δ are required to cover the surface of the \sqrt{n} -sphere? Theorem 5 gives an algorithmic lower bound for this question for every δ ,

Theorem 8. *If m halfspaces of Gaussian measure $\delta < \frac{1}{2}$ cover the \sqrt{n} -sphere in \mathbb{R}^n , then*

$$\delta \geq \Omega\left(\frac{n}{m\sqrt{\log(1 + \frac{m}{n})}}\right).$$

In order to reduce Theorem 7 to Theorem 8, all we require is a restatement of the classical fact attributed to Poincaré (but see [DF87] for a more complete history) that the first coordinate of a uniformly random point on the sphere is approximately distributed like $\mathcal{N}(0, \frac{1}{n})$; modern formulations appear in [Sta82, Spr07]. However, we need a more quantitative version of the bound (Lemma 3.1), and we spend some work showing this.

This approach also explains the assumption of the theorem $\delta \geq 2^{-o(\sqrt{n})}$, a setting which we call the *large cap regime*. When the spherical caps have at least this volume, the approximation between spherical and Gaussian space is good enough to deduce Theorem 7 from Theorem 8. More specifically, if $C = \{x \in S^{n-1} : \langle x, p \rangle \geq \cos \theta\}$ is a spherical cap, then the halfspace $H = \{x \in \mathbb{R}^n : \langle x, p \rangle \geq \sqrt{n} \cos \theta\}$ has $\gamma(H) \approx \text{vol}(C)$ within a constant factor.

An important regime for Conjecture 2 that is outside the large cap regime is when $\theta \in (0, \pi/2)$ is a constant. The corresponding δ is

$$\delta = \Theta\left(\frac{\sin^n \theta}{\sqrt{n}}\right).$$

For example, the case $\theta = \pi/6$ corresponds to the extensively-studied kissing numbers [BDM12, JJP18]. For δ in this range, the approximation in Lemma 3.1 is off by an exponential factor, and one cannot deduce *any* nontrivial spherical measure lower bound.

Note that the known spherical cap constructions show that Theorem 8 is tight up to a log factor provided $\delta \geq 2^{-o(\sqrt{n})}$. When δ is significantly smaller than $2^{-\sqrt{n}}$, Theorem 8 is not tight and the volume bound for spherical caps provides an exponential improvement.

There is a slightly more natural version of the Gaussian covering problem which unfortunately we are not able to resolve. We say that a set S is a $(1 - \varepsilon)$ -cover of a random variable X if $\Pr[X \notin S] \leq \varepsilon$. We say a family of sets is a $(1 - \varepsilon)$ -cover if their union is. If m halfspaces of Gaussian measure δ are a $1/2$ -cover of a standard Gaussian random variable, then there is a trivial volume lower bound $m\delta \geq 1/2$. We conjecture a linear density lower bound for this covering problem,

Conjecture 3. *If m halfspaces of Gaussian measure $\delta < \frac{1}{2}$ are a $1/2$ -cover of a standard n -dimensional Gaussian random variable, then*

$$m \cdot \delta \geq \Omega(n).$$

The choice of $1/2$ corresponds to the fact that the Gaussian measure of the ball of radius \sqrt{n} is approximately $1/2$. Therefore a natural strategy to $1/2$ -cover the Gaussian random variable is to cover the surface of the \sqrt{n} -sphere, for which the bound in Theorem 8 applies. However, we are unable to rule out the existence of $1/2$ -covers with smaller total density.

3.1 Gaussian Space Lower Bound.

We now prove Theorem 8. We need the following fact,

Fact 3.1. *Let φ and $\bar{\Phi}$ denote the PDF and tail probability respectively of the standard normal distribution,*

$$\varphi(t) = \frac{1}{\sqrt{2\pi}} \exp(-t^2/2) \quad \bar{\Phi}(t) = \int_t^\infty \varphi(s) ds$$

There is a constant C so that for $t \geq 1$, $C \frac{\varphi(t)}{t} \leq \bar{\Phi}(t) \leq \frac{\varphi(t)}{t}$.

Proof. See Appendix A.

As can be seen in the proof, the reason for such a tight error bound in Theorem 5 is that the error appears in the exponent here, meaning even small error becomes amplified. Now we prove Theorem 8.

Proof. Let $\{H_i\}$ be a set of m halfspaces each of Gaussian measure δ whose union covers the \sqrt{n} -sphere and assume for the moment that $m \geq 16n$. Let v_i be a normal unit vector to H_i . By Theorem 5 we can find a vector x in the \sqrt{n} -sphere such that

$$\langle v_i, x \rangle \leq \sqrt{2 \log \frac{m}{n}} \cdot \left(1 + O\left(\frac{1}{\log \frac{m}{n}}\right) \right)$$

for each i . Let $\epsilon = O\left(\frac{1}{\log \frac{m}{n}}\right)$ denote the error term.

By assumption x is covered by one of the halfspaces H_i , expressed as

$$H_i = \{x \in \mathbb{R}^n \mid \langle v_i, x \rangle \geq \bar{\Phi}^{-1}(\delta)\}.$$

Therefore,

$$\begin{aligned} \bar{\Phi}^{-1}(\delta) &\leq \sqrt{2 \log \frac{m}{n}} \cdot (1 + \epsilon) \\ \delta &\geq \bar{\Phi}\left(\sqrt{2 \log \frac{m}{n}} \cdot (1 + \epsilon)\right). \end{aligned}$$

Applying Proposition 3.1,

$$\begin{aligned} \delta &\geq C \exp\left(-\left(\sqrt{2 \log \frac{m}{n}} \cdot (1 + \epsilon)\right)^2 / 2 \cdot \frac{(1 + \epsilon)^{-1}}{\sqrt{2 \log \frac{m}{n}}}\right) \\ &= \Theta\left(\left(\frac{n}{m}\right)^{1+O(\epsilon)} \cdot \frac{1}{\sqrt{\log \frac{m}{n}}}\right). \end{aligned}$$

The component of the exponent $\epsilon = O\left(\frac{1}{\log \frac{m}{n}}\right)$ contributes a multiplicative constant factor, and we have the claim,

$$\delta \geq \Omega\left(\frac{n}{m \sqrt{\log(1 + \frac{m}{n})}}\right).$$

We deal with the case in which there are fewer than $16n$ halfplanes. Add vectors until we have $16n$. The theorem in this case tells us

$$\delta \geq \Omega\left(\frac{n}{16n \sqrt{\log \frac{16n}{n}}}\right) = \Omega(1)$$

By the Lusternik-Schnirelmann theorem $m \geq n$, therefore $m\delta \geq \Omega(n)$. □

3.2 From Gaussian Space to the Sphere.

Here we prove Theorem 7 from Theorem 8. Recall our strategy: let C be a spherical cap on S^{n-1} , say with equation

$$C = \{x \in S^{n-1} \mid \langle x, v \rangle \geq \cos(\theta)\}$$

and let H be the halfspace with the same intersection but on the \sqrt{n} -sphere,

$$H = \{x \in \mathbb{R}^n \mid \langle x, v \rangle \geq \cos(\theta)\sqrt{n}\}.$$

We want $\text{vol}(C) \approx \gamma(H)$ to deduce Theorem 7 from Theorem 8. For angle $\theta = \pi/2$, both shapes reduce to hemispheres, and for angle $\theta \geq \pi/2 - O(n^{-1/2})$ both shapes have the

same measure $\Omega(1)$ up to $o(1)$ via the classical theorem that $\sqrt{n}x_1 \xrightarrow{d} N(0, 1)$ for $x \in_{\mathbb{R}} S^n$. The following lemma shows that the two have similar measure for angular radius at least $\pi/2 - o(n^{-1/4})$, or equivalently when one shape has measure at least $2^{-o(\sqrt{n})}$,

Lemma 3.1. *Let $C \subset S^{n-1}$ be a spherical cap with angular radius θ . Let*

$$H = \{x \in \mathbb{R}^n \mid \langle x, v \rangle \geq \sqrt{n} \cos \theta\}$$

where v is some fixed vector, let $\phi = \pi/2 - \theta$, and assume $\phi = o(n^{-1/4})$. Then

$$\text{vol}(C) \sim \gamma(H) \sim \bar{\Phi}(\sqrt{n}\phi).$$

The intuition on why this is the correct assumption on θ is that both C and H have the same intersection with $\{-1, +1\}^n$, and for this θ both $\text{vol}(C)$ and $\gamma(H)$ are well estimated by sampling a uniform boolean point. To explain further, the particular halfspace $H = \{x \in \mathbb{R}^n \mid \langle x, \frac{v}{\sqrt{n}} \rangle \geq t\}$ has Gaussian measure $\bar{\Phi}(t)$. On the other hand, the measure can be estimated by sampling $X \in_{\mathbb{R}} \{+1, -1\}^n$. Setting $S = \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i$, by the Central Limit Theorem we expect $\Pr[S \geq t] \approx \bar{\Phi}(t)$; this is a good approximation for constant t , and in fact for t up to $n^{1/4}$. However, after this point the tail probability becomes exponentially smaller (the subject of large deviation theory is to determine the correct exponent, which in this case is given for all t by the Chernoff-Hoeffding theorem). One interpretation of the lemma is that we prove the boolean sampling procedure also accurately estimates $\text{vol}(C)$ in the range where it estimates $\gamma(H)$, namely $t \leq n^{1/4}$.

There is another geometric interpretation of the lemma. Let G be the cone in \mathbb{R}^n that contains C ,

$$G = \left\{ x \in \mathbb{R}^n : \left\langle \frac{x}{\|x\|_2}, v \right\rangle \geq \cos(\theta) \right\}.$$

Due to the rotational symmetry of the Gaussian measure, $\text{vol}(C) = \gamma(G)$. The halfspace and the cone are two natural bodies that pass through C and $\text{vol}(C) \approx \gamma(H)$ is equivalent to asking that they have similar Gaussian measure.

Proof. The Gaussian measure of H is

$$\begin{aligned} \gamma(H) &= \bar{\Phi}(\sqrt{n} \cos \theta) = \bar{\Phi}(\sqrt{n} \sin \phi) \\ &= \bar{\Phi}(\sqrt{n}\phi - O(\sqrt{n}\phi^3)) \\ &= \bar{\Phi}(\sqrt{n}\phi) + \int_{\sqrt{n}\phi - O(\sqrt{n}\phi^3)}^{\sqrt{n}\phi} \varphi(x) dx \\ &\leq \bar{\Phi}(\sqrt{n}\phi) + O(\sqrt{n}\phi^3) \frac{e^{-(n\phi^2 - O(n\phi^4))/2}}{\sqrt{2\pi}}. \end{aligned}$$

We want to show that the second term is negligible compared to the first. Using the assumption that $\phi = o(n^{-1/4})$,

$$= \bar{\Phi}(\sqrt{n}\phi) + O(\sqrt{n}\phi^3) \frac{e^{-n\phi^2/2}}{\sqrt{2\pi}} (1 + o(1)).$$

We distinguish two cases. When $\phi \geq 1/\sqrt{n}$, by Proposition 3.1 we have

$$\bar{\Phi}(\sqrt{n}\phi) = \Omega\left(\frac{e^{-n\phi^2/2}}{\sqrt{n}\phi}\right).$$

The ratio of the second term to the first is therefore

$$O(n\phi^4) = o(1).$$

In the second case, $\phi < 1/\sqrt{n}$. The first term is at least $\bar{\Phi}(1) = \Omega(1)$, whereas the second term is at most $O(1/n)$. Put together, this shows $\gamma(H) = \bar{\Phi}(\sqrt{n}\phi)(1 + o(1))$.

Now we turn to C . A simple formula for $\text{vol}(C)$ is given in [Li11],

$$\begin{aligned} \text{vol}(C) &= \frac{1}{\sqrt{\pi}} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} \int_0^\theta \sin^{n-2} x \, dx \\ &= \frac{\sqrt{n}}{\sqrt{2\pi}} (1 + o(1)) \int_\phi^{\pi/2} \cos^{n-2} x \, dx \end{aligned}$$

Making the substitution $x = y/\sqrt{n}$,

$$= \frac{1}{\sqrt{2\pi}} (1 + o(1)) \int_{\sqrt{n}\phi}^{\sqrt{n}\pi/2} \cos^{n-2}(y/\sqrt{n}) \, dy.$$

The integrand is approximately $\cos^n(y/\sqrt{n}) \approx (1 - y^2/2n)^n \approx e^{-y^2/2}$, so at least heuristically we have the claim $\text{vol}(C) \approx \bar{\Phi}(\sqrt{n}\phi)$. To make this argument rigorous, let $T = \max(n^{0.2}, n^{0.7}\phi)$ and break the integral into two pieces,

$$\begin{aligned} \frac{1}{\sqrt{2\pi}} \int_{\sqrt{n}\phi}^{\sqrt{n}\pi/2} \cos^{n-2}(y/\sqrt{n}) \, dy &= \frac{1}{\sqrt{2\pi}} \int_{\sqrt{n}\phi}^T \cos^{n-2}(y/\sqrt{n}) \, dy \\ &\quad + \frac{1}{\sqrt{2\pi}} \int_T^{\sqrt{n}\pi/2} \cos^{n-2}(y/\sqrt{n}) \, dy. \end{aligned}$$

The first piece can be Taylor expanded as the upper bound of integration is $o(\sqrt{n})$. The integrand is

$$\begin{aligned} \cos^{n-2}\left(\frac{y}{\sqrt{n}}\right) &= \left(1 - \frac{y^2}{2n} + O\left(\frac{y^4}{n^2}\right)\right)^{n-2} \\ &= \left(1 - \frac{y^2}{2n}\right)^n \left(1 + O\left(\frac{y^4}{n^2}\right)\right)^n \cdot (1 + o(1)) \end{aligned}$$

Proposition A.1, in the appendix, shows that the first term is asymptotic to $e^{-y^2/2}$ and the second term is $1 + o(1)$, with both using the assumption that $y \leq O(\sqrt{n}\phi) = o(n^{1/4})$. The first piece is therefore equal to

$$\frac{1 + o(1)}{\sqrt{2\pi}} \int_{\sqrt{n}\phi}^T e^{-y^2/2} \, dy.$$

The second piece is exponentially smaller than the first. However, we leave this claim unproven for a moment in order to show that the first piece is $(1 + o(1))\bar{\Phi}(\sqrt{n}\phi)$. That is, we want to show that

$$\int_{\sqrt{n}\phi}^T e^{-y^2/2} dy \geq (1 - o(1)) \int_{\sqrt{n}\phi}^{\infty} e^{-y^2/2} dy.$$

The tail is upper bounded by Proposition 3.1,

$$\int_T^{\infty} e^{-y^2/2} dy \leq \frac{e^{-T^2/2}}{T}.$$

On the other hand Proposition 3.1 gives us a lower bound on $\bar{\Phi}(\sqrt{n}\phi)$, at least in the case when $\phi \geq 1/\sqrt{n}$,

$$\bar{\Phi}(\sqrt{n}\phi) \geq C \frac{e^{-n\phi^2/2}}{\sqrt{n}\phi}.$$

The exponent $e^{-T^2/2} = O(e^{-n^{1.4}\phi^2/2})$ is exponentially smaller. If $\phi < 1/\sqrt{n}$, then $\bar{\Phi}(\sqrt{n}\phi) = \Omega(1)$, whereas the tail is bounded by $1/T = O(n^{-0.2})$. In either case, the tail is $o(\bar{\Phi}(\sqrt{n}\phi))$ as needed.

Now we return to bounding the second piece. We show that it's negligible compared to $\bar{\Phi}(\sqrt{n}\phi)$. Bounding the integrand by the left endpoint shows it's at most $\frac{\sqrt{n}\pi}{2} \cdot \cos^{n-2}(T)$ which as done above can be approximated using the Taylor expansion,

$$= (1 + o(1)) \cdot \frac{\sqrt{n}\pi}{2} \cdot e^{-T^2/2}.$$

In the case that $\phi \geq 1/\sqrt{n}$, this is $O(e^{-n^{1.4}\phi^2})$, whereas by Proposition 3.1, $\bar{\Phi}(\sqrt{n}\phi)$ has exponent $e^{-n\phi^2}$, and hence this term is exponentially smaller. In the case that $\phi < 1/\sqrt{n}$, this is $O(\sqrt{n}e^{-n^{0.4}/2})$ whereas $\bar{\Phi}(\sqrt{n}\phi) = \Omega(1)$. This shows that the second piece of the integral is negligible, while the first is asymptotic to $\bar{\Phi}(\sqrt{n}\phi)$, which completes the proof of the lemma. \square

Finally, we fill in the details of the proof of Theorem 7 using Theorem 8.

Proof. Say we have a collection of m caps $\{C_i\}$ whose union covers S^{n-1} , and each has angular radius θ and normalized volume δ with $2^{-o(\sqrt{n})} \leq \delta < \frac{1}{2}$. Let v_i be the pole of cap C_i , and define the halfspaces

$$H_i = \{x \in \mathbb{R}^n \mid \langle x, v_i \rangle \geq \sqrt{n} \cos \theta\}.$$

The intersection of H_i with the \sqrt{n} -sphere is exactly $\sqrt{n}C_i$, and the assumption that the $\{C_i\}$ cover S^{n-1} tells us that the H_i cover the \sqrt{n} -sphere. Apply Theorem 8 to the collection of $\{H_i\}$,

$$\gamma(H_i) \geq \Omega\left(\frac{n}{m\sqrt{\log(1 + \frac{m}{n})}}\right).$$

We would like to now apply Lemma 3.1; to do so we need to bound θ given δ . By Proposition A.2, a halfspace of Gaussian volume δ has

$$\bar{\Phi}^{-1}(\delta) = \sqrt{2 \ln 1/\delta} + o(1) = o(n^{1/4}).$$

Using Lemma 3.1, the spherical cap through this halfspace has measure $\delta(1 + o(1))$ and angle $\frac{\pi}{2} - o(n^{-1/4})$. Enlarging the angle by $o(n^{-1/4})$ is enough to contain the spherical cap of volume δ .

Now by Lemma 3.1, $\gamma(H_i)$ is within a constant factor of δ , and therefore

$$\delta \geq \Omega \left(\frac{n}{m \sqrt{\log(1 + \frac{m}{n})}} \right).$$

□

3.3 A Matching Bound for Theorem 5.

The lemmas established in this section can be used to convert the cap covering bounds of Böröczky and Wintsche [BW03] to bounds for Theorem 5. We don't make too much fuss about the constants for covers and suffice to work with the simpler covering density bound of $2n \ln n$.

Theorem 6. *For every choice of $2^{-o(\sqrt{n})} < \delta < \frac{1}{n^2}$, there is a parameter $1/\delta \leq m \leq \frac{2n \ln n}{\delta}$ and unit vectors $v_1, \dots, v_m \in S^{n-1}$ such that, for any $x \in S^{n-1}$ there is a v_i with*

$$\langle x, v_i \rangle \geq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 - o \left(\frac{1}{\sqrt{\log \frac{m}{n}}} \right) \right).$$

Proof. Fix δ . Find a set of caps of m caps of volume δ with polar vectors v_i that cover the sphere, with

$$1 \leq m\delta \leq 2n \ln n.$$

Letting θ be the angular radius of the caps, for any $x \in S^{n-1}$ there is v_i with $\langle v_i, x \rangle \geq \cos \theta$. The calculation performed in the proof of Theorem 7 showed that θ was in the correct range for Lemma 3.1.

$$\begin{aligned} \bar{\Phi}(\sqrt{n} \cos \theta) &= \delta(1 + o(1)) \\ \cos \theta &= \frac{1}{\sqrt{n}} \bar{\Phi}^{-1}(\delta(1 + o(1))) \end{aligned}$$

Taylor expand $\bar{\Phi}^{-1}$ via Proposition A.2, given in the appendix,

$$\begin{aligned} \bar{\Phi}^{-1}(\delta(1 + o(1))) &= \sqrt{2 \ln 1/\delta} + o(1) \\ &\geq \sqrt{2 \ln \frac{m}{n} - 2 \ln(2 \ln n)} + o(1) \end{aligned}$$

$$= \sqrt{2 \ln \frac{m}{n}} + o(1).$$

Put together, we have

$$\langle v_i, x \rangle \geq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \left(1 + o\left(\frac{1}{\sqrt{\log \frac{m}{n}}}\right) \right).$$

□

3.4 Generating Sphere Packings.

We verify the performance of Algorithm 2.1 for generating sphere packings,

Theorem 9. *Let $16n \leq m \leq 2^{o(\sqrt{n})}$. There is a deterministic algorithm which runs in time $\text{poly}(m, n)$ that outputs m points $v_1, v_2, \dots, v_m \in S^{n-1}$ such that*

$$\langle v_i, v_j \rangle \leq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \cdot \left(1 + O\left(\frac{1}{\log \frac{m}{n}}\right) \right)$$

for all pairs i, j . Taking the maximal radius r such that spherical caps of radius r around the v_i are disjoint produces a packing with density $\Omega\left(\frac{n}{2^n \sqrt{\log \frac{m}{n}}}\right)$.

The density of this packing density upper bound is related to the density of the covering lower bound by a factor of 2^n . This is (not directly due to, but essentially) because of the well-known generic relationship between packing and covering: doubling the radius of a maximal packing produces a covering, with density scaled by 2^n . However, this generic relationship is in general not tight.

Proof. The first part is immediate from Theorem 5. For the second part, let $2r$ be the minimum distance between the v_i . Taking a cap of radius $2r$ around each point covers S^{n-1} , and the density of this covering was computed in Theorem 7 to be at least $\Omega\left(\frac{n}{\sqrt{\log \frac{m}{n}}}\right)$. Halving the radius gives a disjoint set of caps of the desired density. □

4 Komlós Problem in the Spherical Domain.

Seeing as the spherical domain sidesteps the rounding present in partial coloring, which seems to have inherent flaws when trying to prove tight discrepancy results in the boolean domain (see [BG17] for some discussion), it is natural to hope that the spherical versions of some open problems from discrepancy theory are more tractable. In particular we can adapt our Algorithm 2.1 to resolve a version of the Komlós problem in the spherical setting,

Theorem 12. Let $w_1, \dots, w_n \in \mathbb{R}^n$ be vectors with $\|w_i\|_2 \leq 1$, and let W be the matrix with the w_i as columns. Then we can find a unit vector $x \in S^{n-1}$ such that

$$\|Wx\|_\infty = O\left(\frac{1}{\sqrt{n}}\right)$$

in time polynomial in n .

The idea is a rather general one taken from boolean discrepancy theory. A priori some of the rows of W could have norm as large as \sqrt{n} . Even so, in the special case of the Komlós conjecture where the rows of W are restricted to have norm $O(1)$, the core of the problem still seems to be captured; for many algorithms for discrepancy problems, the row norms are assumed to be bounded without loss of generality. This holds for algorithms that iteratively build a solution from small update vectors subject to linear number linear number γn ($\gamma < 1$) of linear constraints. The bounds $\|w\|_2 \leq 1$ ensure there are at most $\varepsilon^2 n$ row vectors with norm greater than $1/\varepsilon$. Thus as long as $\gamma + \varepsilon^2 < 1$, the number of linear constraints we have is significantly less than n , and there is still room to find an update vector orthogonal to rows of large norm.

In fact, as pointed out by the anonymous reviewer, the above idea can be instantiated in a black box manner using just the generic partial coloring lemma of Lovett and Meka,

Theorem 13 [LM15]. Let $v_1, \dots, v_m \in \mathbb{R}^n$ be vectors, $x^{(0)} \in [-1, +1]^n$ be a starting point, and let $c_1, \dots, c_m \geq 0$ be thresholds so that $\sum_{j=1}^m \exp(-c_j^2/16) \leq \frac{n}{32}$. Let $\delta > 0$ be a small approximation parameter. Then there exists a $\text{poly}(m, n, 1/\delta)$ -time randomized algorithm which with probability at least 0.1 finds a point $x \in [-1, +1]^n$ such that

- (i) $|\langle v_j, x - x^{(0)} \rangle| \leq c_j \|v_j\|_2$ for all $j \in [m]$,
- (ii) $|x_i| \geq 1 - \delta$ for at least $n/2$ indices $i \in [n]$.

We now give the details for Theorem 12.

Proof. We will invoke Theorem 13. Let W be the given matrix, and let the v_j be the rows of W . Let B be the set of v_j with norm at least $\|v_j\| \geq 20$. Set the starting point $x^{(0)} = 0$ and $\delta = 0.5$, and set the parameters

$$c_j = \begin{cases} 0 & j \in B \\ 50 & j \notin B \end{cases}$$

We must check $\sum_j \exp(-c_j^2/16) \leq \frac{n}{32}$. Because of the unit norm constraints on the columns of W , the set B isn't that big. The sum $\sum_j \|v_j\|^2$ is at most n , and hence at most $n/400$ of the v_j have $\|v_j\| \geq 20$. Thus we can bound

$$\begin{aligned} \sum_j \exp(-c_j^2/16) &= \sum_{j \in B} \exp(-c_j^2/16) + \sum_{j \notin B} \exp(-c_j^2/16) \\ &\leq \frac{n}{400} + n \cdot \exp(-50^2/16) \leq \frac{n}{32}. \end{aligned}$$

Invoking the theorem, we can algorithmically produce x . We will show $\hat{x} = \frac{x}{\|x\|}$ satisfies our claim. Because at least half of the coordinates of x are $+1$ or -1 , we have $\|x\| = \Omega(\sqrt{n})$. For rows in B , the inner product between x and v_j is exactly zero, and hence these rows are definitely $O(1/\sqrt{n})$. For rows outside of B , the norm of v_j is $O(1)$ and from the theorem we have $\langle x, v_j \rangle \leq O(1) \cdot \|v_j\| = O(1)$. Therefore, the inner product with the normalized vector \hat{x} is $O(1/\sqrt{n})$ as desired. \square

5 Hardness of Approximation.

In this section we show that the SPHERICAL DISCREPANCY problem is APX-hard. Here is the formal specification of the SPHERICAL DISCREPANCY problem,

SPHERICAL DISCREPANCY

Input: a collection of unit vectors $v_1, v_2, \dots, v_m \in S^{n-1}$

Output: compute the minimum value of $\max_i \langle x, v_i \rangle$ for $x \in S^{n-1}$

We say that an algorithm is an α -approximation or α -approximates SPHERICAL DISCREPANCY if it outputs $x \in S^{n-1}$ which is within a multiplicative α factor of the best possible x ,

$$\max_i \langle v_i, x \rangle \leq \alpha \cdot \min_{x \in S^{n-1}} \max_i \langle v_i, x \rangle$$

We prove a constant factor hardness of distinguishing result for SPHERICAL DISCREPANCY,

Theorem 11. *There is a constant $C > 1$ so that it is NP-hard to distinguish between instances of SPHERICAL DISCREPANCY with $m = O(n)$ with value at most $\frac{1}{\sqrt{n}}$, and instances with value at least $\frac{C}{\sqrt{n}}$.*

Corollary 1. *For some constant $C > 1$, it is NP-hard to C -approximate SPHERICAL DISCREPANCY.*

Previous work by Petković et al [PPL12] using a different approach showed that solving SPHERICAL DISCREPANCY exactly is NP-hard when the v_i are not restricted to be unit vectors.

A few words are in order about the notion of approximation algorithm. For boolean and spherical discrepancy the notion of α -approximation is slightly orthogonal to the goals of papers such as [Spe85, BDGL18]. This is because for many discrepancy problems, we don't even know the worst-case value of the optimum solution across the possible inputs,

$$\max_{v_1, v_2, \dots, v_m} \min_{x \in S^{n-1}} \max_i \langle x_i, v \rangle.$$

Theorems such as Spencer's or the Komlós conjecture are attempts to prove (or, in the case of an algorithm, algorithmically certify) an upper bound on the above quantity; however, they often provide no guarantee of α -approximation because they do not relate the found algorithmic solution to the value of an optimum solution for the input. In some ways this

is the “interesting” challenge for discrepancy because the problem is much harder when the algorithm is forced to make a guarantee that depends on the optimum solution: for example it is NP-hard to distinguish boolean set systems that have discrepancy zero from those with discrepancy $\Omega(\sqrt{n})$ [CNN11]. We conjecture in Section 6 that for SPHERICAL DISCREPANCY the situation is similar and the true approximation factor is significantly worse than what is proven here.

The reduction used to prove Theorem 11 is a gap-preserving gadget reduction from MAX NAE-E3-SAT. In the MAX NAE-E3-SAT problem, which stands for Not-All-Equal Exactly-3 SAT, we are given a collection of m clauses each of which involves exactly three distinct literals. An assignment to the variables satisfies a clause if the assignments to all literals are not all the same. The instances must also have the number of occurrences of each variable bounded by some absolute constant B . An observation made by Charikar et al [CGW05, Theorem 11] states that (even further restricting the SAT instance to be monotone) there are constants $\gamma < 1$ and B so that it is NP-hard to distinguish an instance which is satisfiable, from one in which at most γm clauses can be simultaneously satisfied. Observe that the size of these instances is guaranteed to be linear, $m \leq \frac{Bn}{3} = O(n)$, by construction. Now we prove Theorem 11.

Proof. Let C_1, \dots, C_m be a hard instance of MAX NAE-E3-SAT. The dimension for our SPHERICAL DISCREPANCY instance will be n . Construct the instance as follows:

- For each C_i take its $\{0, \pm 1\}$ indicator vector $\mathbb{1}_{C_i}$ for whether a variable occurs in the clause and whether the variable is negated in the clause. Add the two vectors $\frac{1}{\sqrt{3}}\mathbb{1}_{C_i}$ and $\frac{-1}{\sqrt{3}}\mathbb{1}_{C_i}$.
- Add the vectors e_i and $-e_i$ for $i = 1, 2, \dots, n$.

It is clear that all vectors are unit vectors and the number of vectors is $O(n)$. If the MAX NAE-E3-SAT instance is satisfiable, let x be the normalized $\{\pm 1\}$ -coloring. On vectors of the first type, the NAE constraint implies $|\langle x, v_i \rangle| = \frac{1}{\sqrt{3n}}$, whereas on the second type $|\langle x, v_i \rangle| = \frac{1}{\sqrt{n}}$, so the value of the new instance is $\frac{1}{\sqrt{n}}$.

On the other hand, assume the MAX NAE-E3-SAT instance is far from satisfiable, and let x be a unit vector. We must show the value of x is at least $\frac{C}{\sqrt{n}}$ for some constant $C > 1$. On boolean coordinates $x \in \{\frac{1}{\sqrt{n}}, \frac{-1}{\sqrt{n}}\}^n$, since there is at least one clause which is not satisfied by the corresponding assignment, the value of x is exactly $\sqrt{\frac{3}{n}}$, which is larger than $\frac{1}{\sqrt{n}}$, hence the value is large here. But as x moves away from the set $\{\frac{1}{\sqrt{n}}, \frac{-1}{\sqrt{n}}\}^n$, the value is forced to be large because of the vectors of the second type. We convert this intuition to a proof.

Choose $C = \min \left(\sqrt{\frac{B - \frac{9}{16}(1-\gamma)}{B - (1-\gamma)}}, \frac{3\sqrt{3}}{4} \right)$. If $\|x\|_\infty \geq \frac{C}{\sqrt{n}}$, the value of this solution will be at least $\frac{C}{\sqrt{n}}$ via the vectors of the second type, so assume that $\|x\|_\infty < \frac{C}{\sqrt{n}}$.

Let $S = \{i : x_i < \frac{0.75}{\sqrt{n}}\}$, a set we call the “small” coordinates of x , and let $\alpha = |S|/n$. Let \tilde{x} denote the coloring obtained by taking the sign of each coordinate of x (assign arbitrarily

for zero). There are at least $1 - \gamma$ fraction of clauses not satisfied by \tilde{x} . We claim that for some clause C_i not satisfied by \tilde{x} , none of its three variables are in S . This will finish the proof: because the coordinates of x must have the same signs in C_i , and they are all big, the inner product with $\frac{1}{\sqrt{3}}\mathbb{1}_{C_i}$ or $\frac{-1}{\sqrt{3}}\mathbb{1}_{C_i}$ must be least $3 \cdot \frac{.75}{\sqrt{3n}} \geq \frac{C}{\sqrt{n}}$. Suppose then that every unsatisfied clause by \tilde{x} has a variable from S . Because of the bound B on the maximum number of occurrences of any variable, $\alpha \geq \frac{1-\gamma}{B}$. Outside of S we use the bound on $\|x\|_\infty$ to bound the big coordinates,

$$\begin{aligned} \sum_{i=1}^n x_i^2 &< \left(\frac{C}{\sqrt{n}}\right)^2 (1 - \alpha) \cdot n + \left(\frac{0.75}{\sqrt{n}}\right)^2 \alpha \cdot n \\ &\leq C^2 \left(1 - \frac{1-\gamma}{B}\right) + \frac{9}{16} \cdot \frac{1-\gamma}{B} \end{aligned}$$

The choice of $C \leq \sqrt{\frac{B - \frac{9}{16}(1-\gamma)}{B - (1-\gamma)}}$ shows the right-hand side is at most 1, contradicting that $\sum_i x_i^2 = 1$. \square

6 Further Questions.

The current work implements an algorithmic approach to covering problems in geometry, and there are several questions left open for future investigation.

- Perhaps most generally, the idea to generalize $\{-1, +1\}^n$ to the \sqrt{n} -radius sphere may be interesting to consider for other combinatorial problems besides discrepancy. We would be interested in seeing concrete realizations of the following algorithmic strategy:
 - (1) Relax an optimization problem over $\{-1, +1\}^n$ to the \sqrt{n} -radius sphere.
 - (2) Solve the relaxed problem.
 - (3) Round the \sqrt{n} -radius sphere to $\{-1, +1\}^n$.

For example, Trevisan [Tre12] provides an interesting algorithm for MaxCut that exactly matches this framework. Indeed, spectral relaxations such as Trevisan’s can be interpreted in this perspective. However, in general, unlike spectral algorithms, in our framework it is not immediate that the relaxed problem in step (2) can be perfectly solved! Spectral relaxation corresponds to objective functions of degree at most 2 which can be perfectly solved, but polynomial optimization of higher degree over the sphere is NP-hard in general (and even difficult to approximate [BGL⁺17]). Therefore we find it particularly interesting whether any problems which are naturally degree-4 or higher (or, as in the case of discrepancy, are not polynomials) can be attacked using this framework. The hypersphere also provides more structure for rounding than e.g. a Sum-of-Squares relaxation.

- If we were able to improve the error bound in Algorithm 2.1, we could remove the logarithmic factor in Theorem 7. The conjectural “right” error bound for proving a linear lower bound can be computed from Proposition A.2 with $\delta = \frac{n}{m}$,

Conjecture 4. *There is an efficient algorithm that improves the bound in Theorem 5 to*

$$\langle v_i, x \rangle \leq \sqrt{\frac{2 \ln \frac{m}{n}}{n}} \left(1 - \frac{\ln \ln \frac{m}{n}}{4 \ln \frac{m}{n}} + O\left(\frac{1}{\ln \frac{m}{n}}\right) \right).$$

- However, even if we are able to remove the log factor, it is unclear if our techniques are able to prove the bona fide simplex bound $\tau_{n,\theta}$, instead of the bound $\Omega(n)$, which is weaker by a constant factor. It seems possible that the essence of Algorithm 2.1 could be extracted into a continuous-time walk (either deterministic or random) which avoids the lossy algorithmic analysis we went through.

Question 1. *Can Algorithm 2.1 be refined to improve the constant in the lower bound and completely prove Conjecture 2?*

- We should point out some fundamental geometric questions about $\tau_{n,\delta}$ that, to the best of the authors' knowledge, are open. The formal definition of $\tau_{n,\delta}$ is as follows. Let T be a regular spherical simplex inscribed in a spherical cap of volume δ . Let C_i be caps of volume δ around vertex i of T . Then

$$\tau_{n,\delta} \stackrel{\text{def}}{=} \frac{\sum_{i=1}^n \text{vol}(C_i \cap T)}{\text{vol}(T)} = \frac{n \cdot \text{vol}(C_1 \cap T)}{\text{vol}(T)}.$$

Rogers [Rog58] computed that for $\tau_n = \lim_{\delta \rightarrow 0} \tau_{n,\delta}$ the Euclidean covering density, $\tau_n \sim \frac{n}{e\sqrt{e}}$. It is natural to conjecture that the densities are always linear,

Conjecture 5. $\tau_{n,\delta} \geq c \cdot n$ for some positive constant c .

The conjecture is verified in the “nearly Euclidean” regime by Böröczky and Wintsche [BW03, Example 6.3]. Intuitively, as the simplex becomes more curved, the relative volume near the center of the simplex increases. The caps C contain approximately half of the local mass near the center, and so we might expect the cap to contain more and more of the simplex as it becomes more curved, up to the limit where T is a hemisphere and each cap equals half of the simplex, $\tau_{n,1/2} = n/2$. Based on this intuition we conjecture monotonicity of $\tau_{n,\delta}$,

Conjecture 6. *For every n , as a function of the dihedral angle $\theta \in [\arccos(1/(n-1)), \pi]$ the function $\tau_{n,\theta}$ is monotonically increasing.*

Here we have changed notation, so that θ is the angle between two planes defining the simplex, and we have extended the conjecture to include hyperbolic space (the minimum dihedral angle in n -dimensional hyperbolic space is $\arccos(1/(n-1))$, achieved by the ideal regular simplex).

One must be careful, however, because the related expression for the simplicial packing density is a *decreasing* function of θ (for sufficiently large n), as proven by Marshall [Mar99] and Kellerhals [Kel98] (though we could not verify the proof outside of hyperbolic space). The proof is analytic, and Kellerhals poses as an open problem to find a geometric proof, which we also pose as a challenge for Conjectures 5 and 6.

- The section on sphere covering presented Conjecture 3,

Conjecture 3. *If m halfspaces of Gaussian measure $\delta < \frac{1}{2}$ are a $1/2$ -cover of a standard n -dimensional Gaussian random variable, then*

$$m \cdot \delta \geq \Omega(n).$$

It may be possible to adapt the boolean discrepancy “sample and project” algorithm of Rothvoss [Rot17] or Eldan and Singh [ES18] to prove this lower bound.

- We briefly pointed out in Theorem 9 that Algorithm 2.1 can also be used to generate a set of points in S^{n-1} that are relatively spaced out. Therefore we ask,

Question 2. *Can we use Algorithm 2.1 to deterministically build smaller hitting sets than Rabani-Shpilka [RS09]? For every $\delta \geq 2^{-o(\sqrt{n})}$, can we deterministically generate a sphere cover using spherical caps of volume δ with density $\tilde{O}(n)$ in time $\text{poly}(n, 1/\delta)$?*

The problem with using Algorithm 2.1 in its current form is that Theorem 5 is a “packing property”; we have no guarantee that m points will cover the whole sphere.

- Algorithm 2.1 minimizes the max of a collection of linear functions on the sphere. Can we adapt it to minimize functions that are “slightly nonlinear”, or sets with boundaries that are slightly nonlinear? An interesting question arises if we consider sets with a diameter bound.

Fix a parameter $\theta \in [0, \pi]$, and define a distance graph in spherical space $S_{\geq \theta}^n$ with vertex set S^n and edge set

$$E(S_{\geq \theta}^n) \stackrel{\text{def}}{=} \{(x, y) \mid \langle x, y \rangle \geq \theta\}.$$

Independent sets in $S_{\geq \theta}^n$ are sets with diameter at most θ , and therefore a cover of S^n by spherical caps with diameter θ yields a finite coloring of $S_{\geq \theta}^n$. It is hopeful that this is the best possible up to a constant,

Conjecture 7. *For every θ , $\chi(S_{\geq \theta}^n) \geq \Omega(B_{n,\theta})$, where $B_{n,\theta}$ equals the minimum number of spheres of radius θ needed to cover S^n .*

- The basic problem of SPHERICAL DISCREPANCY has constant factor hardness of approximation as we showed in Section 5 but it seems likely that the problem has a much worse approximation factor.

Conjecture 8. *For every $16n \leq m \leq 2\sqrt{n}$, it is NP-hard to approximate SPHERICAL DISCREPANCY on m unit vectors within a factor of $\Omega(\sqrt{n \ln \frac{m}{n}})$.*

Evidence for this conjecture comes from the boolean regime, where despite the fact that every set system on $O(n)$ sets has discrepancy $O(\sqrt{n})$, it is NP-hard to distinguish set systems with zero discrepancy from those with discrepancy $\Omega(\sqrt{n})$ [CNN11]. In the spherical domain, on the other hand, given a set of unit vectors it is easy to check if there is a discrepancy zero vector i.e. a vector orthogonal to the entire set. It is not clear how the SPHERICAL DISCREPANCY problem behaves when we promise a lower bound on the solution value such as $\frac{1}{n}$ in order to avoid issues of this sort.

7 Acknowledgments

We thank Will Perkins for clarifying the extent to which the proofs in [CFR59] apply to spherical cap lower bounds, which inspired us to pursue this work further. We thank Andy Drucker for many discussions and guidance on this work, and for introducing us to discrepancy theory.

References

- [AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8:121–164, 2012.
- [Ban12] Nikhil Bansal. Semidefinite optimization in discrepancy theory. *Math. Program.*, 134(1, Ser. B):5–22, 2012.
- [BDG19] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An algorithm for Komlós conjecture matching Banaszczyk’s bound. *SIAM J. Comput.*, 48(2):534–553, 2019.
- [BDGL18] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt walk: a cure for the Banaszczyk blues. In *Proceedings of the 50th Annual Symposium on Theory of Computing*, pages 587–597. 2018.
- [BDM12] Peter Boyvalenkov, Stefan Dodunekov, and Oleg Musin. A survey on the kissing numbers. *Serdica Math. J.*, 38(4):507–522, 2012.
- [BEJ76] J. M. Blair, C. A. Edwards, and J. H. Johnson. Rational Chebyshev approximations for the inverse of the error function. *Math. Comp.*, 30(136):827–830, 1976.
- [BG17] Nikhil Bansal and Shashwat Garg. Algorithmic discrepancy beyond partial coloring. In *Proceedings of the 49th Annual Symposium on Theory of Computing*, pages 914–926. 2017.
- [BGL⁺17] Vijay Bhattiprolu, Venkatesan Guruswami, Euiwoong Lee, Mrinalkanti Ghosh, and Madhur Tulsiani. Weak decoupling, polynomial folds, and approximate optimization over the sphere. In *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, pages 1008–1019. IEEE Computer Soc., Los Alamitos, CA, 2017.
- [Bör78] Károly Böröczky, Jr. Packing of spheres in spaces of constant curvature. *Acta Math. Acad. Sci. Hungar.*, 32(3-4):243–261, 1978.
- [Bör04] Károly Böröczky, Jr. Finite packing and covering. 154:xviii+380, 2004.
- [BW03] Károly Böröczky, Jr. and Gergely Wintsche. Covering the sphere by equal spherical balls. In *Discrete and Computational Geometry: the Goodman-Pollack Festschrift*, pages 235–251. Springer, 2003.

- [CFR59] H. S. M. Coxeter, L. Few, and C. A. Rogers. Covering space with equal spheres. *Mathematika*, 6:147–157, 1959.
- [CGW05] Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. *J. Comput. System Sci.*, 71(3):360–383, 2005.
- [Cha00] Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, Cambridge, 2000.
- [CL09] Frédéric Cazals and Sébastien Lorient. Computing the arrangement of circles on a sphere, with applications in structural biology. *Computational Geometry*, 42(6-7):551–565, 2009.
- [CNN11] Moses Charikar, Alantha Newman, and Aleksandar Nikolov. Tight hardness results for minimizing discrepancy. In *Proceedings of the Twenty-second Annual Symposium on Discrete Algorithms*, pages 1607–1614, 2011.
- [Cox63] H. S. M. Coxeter. An upper bound for the number of equal nonoverlapping spheres that can touch another of the same size. In *Proc. Sympos. Pure Math., Vol. VII*, pages 53–71. Amer. Math. Soc., Providence, R.I., 1963.
- [DF87] Persi Diaconis and David Freedman. A dozen de Finetti-style results in search of a theory. *Ann. Inst. H. Poincaré Probab. Statist.*, 23(2, suppl.):397–423, 1987.
- [ES18] Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *Random Structures Algorithms*, 53(2):289–307, 2018.
- [JJP18] Matthew Jenssen, Felix Joos, and Will Perkins. On kissing numbers and spherical codes in high dimensions. *Adv. Math.*, 335:307–321, 2018.
- [Kel98] Ruth Kellerhals. Ball packings in spaces of constant curvature and the simplicial density function. *J. Reine Angew. Math.*, 494:189–203, 1998. Dedicated to Martin Kneser on the occasion of his 70th birthday.
- [Li11] S. Li. Concise formulas for the area and volume of a hyperspherical cap. *Asian J. Math. Stat.*, 4(1):66–70, 2011.
- [LM15] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.
- [Lov00] László Lovász. Integer sequences and semidefinite programming. *Publ. Math. Debrecen*, 56(3-4):475–479, 2000. Dedicated to Professor Kálmán Györy on the occasion of his 60th birthday.
- [LRR17] Avi Levy, Harishchandra Ramadas, and Thomas Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *Integer programming and combinatorial optimization*, pages 380–391. 2017.
- [Mar99] T. H. Marshall. Asymptotic volume formulae and hyperbolic ball packing. *Ann. Acad. Sci. Fenn. Math.*, 24(1):31–43, 1999.

- [Nik13] Aleksandar Nikolov. The Komlós conjecture holds for vector colorings. E-print, arXiv:1301.4039, 2013.
- [PPL12] Marko D. Petković, Dragoljub Pokrajac, and Longin Jan Latecki. Spherical coverage verification. *Appl. Math. Comput.*, 218(19):9699–9715, 2012.
- [Rog58] C. A. Rogers. The packing of equal spheres. *Proc. London Math. Soc. (3)*, 8:609–620, 1958.
- [Rot17] Thomas Rothvoss. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.*, 46(1):224–234, 2017.
- [RS09] Yuval Rabani and Amir Shpilka. Explicit construction of a small epsilon-net for linear threshold functions. In *Proceedings of the 41st Annual Symposium on Theory of Computing*, pages 649–658. 2009.
- [Spe85] Joel Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289(2):679–706, 1985.
- [Spr07] M. C. Spruill. Asymptotic distribution of coordinates on high dimensional spheres. *Electron. Comm. Probab.*, 12:234–247, 2007.
- [Sta82] A. J. Stam. Limit theorems for uniform distributions on spheres in high-dimensional Euclidean spaces. *J. Appl. Probab.*, 19(1):221–228, 1982.
- [SV05] Achill Schürmann and Frank Vallentin. Local covering optimality of lattices: Leech lattice versus root lattice E_8 . *Int. Math. Res. Not.*, (32):1937–1955, 2005.
- [Tre12] Luca Trevisan. Max cut and the smallest eigenvalue. *SIAM J. Comput.*, 41(6):1769–1786, 2012.

Appendix A

Fact 2.1. For any $0 \leq x \leq 1$,

$$e^x \leq 1 + x + \frac{x^2}{2} + \frac{x^3}{2}.$$

Proof.

$$\begin{aligned} e^x &= 1 + x + \frac{x^2}{2} + \sum_{k=3}^{\infty} \frac{x^k}{k!} \\ &\leq 1 + x + \frac{x^2}{2} + x^3 \left(\sum_{k=3}^{\infty} \frac{1}{k!} \right) \\ &= 1 + x + \frac{x^2}{2} + x^3 (e - 2.5) \\ &\leq 1 + x + \frac{x^2}{2} + \frac{x^3}{2} \end{aligned}$$

□

Fact 3.1. Let φ and $\bar{\Phi}$ denote the PDF and tail probability respectively of the standard normal distribution,

$$\varphi(t) = \frac{1}{\sqrt{2\pi}} \exp(-t^2/2) \quad \bar{\Phi}(t) = \int_t^\infty \varphi(s) ds$$

There is a constant C so that for $t \geq 1$, $C \frac{\varphi(t)}{t} \leq \bar{\Phi}(t) \leq \frac{\varphi(t)}{t}$.

Proof.

$$\bar{\Phi}(t) = \frac{1}{2\pi} \int_t^\infty e^{-x^2/2} dx = \varphi(t) \int_0^\infty e^{-x^2/2} \cdot e^{-xt} dx$$

For the upper bound,

$$\bar{\Phi}(t) \leq \varphi(t) \int_0^\infty e^{-xt} dx = \frac{\varphi(t)}{t}.$$

For the lower bound, since φ decreases so quickly the integral is well-approximated by just the first unit interval,

$$\bar{\Phi}(t) \geq \varphi(t) \int_0^1 e^{-1/2} e^{-xt} dx = \frac{\varphi(t)}{t} \cdot e^{-1/2} (1 - e^{-t}).$$

Assuming that $t \geq 1$ ensures $C = e^{-1/2}(1 - 1/e)$ works. □

Fact A.1. For $x = o(\sqrt{n})$,

$$\left(1 + \frac{x}{n}\right)^n \sim e^x$$

For $x = \Theta(\sqrt{n})$,

$$\left(1 + \frac{x}{n}\right)^n = \Theta(e^x)$$

For $x = \omega(\sqrt{n})$,

$$\left(1 + \frac{x}{n}\right)^n = \omega(e^x)$$

Proof.

$$\begin{aligned} \left(1 + \frac{x}{n}\right)^n &= \exp\left(\ln\left(\left(1 + \frac{x}{n}\right)^n\right)\right) \\ &= \exp\left(n \cdot \ln\left(1 + \frac{x}{n}\right)\right) \\ &= \exp\left(n \cdot \left(\frac{x}{n} + \Theta\left(\frac{x^2}{n^2}\right)\right)\right) \\ &= e^x \cdot \exp(\Theta(x^2/n)) \end{aligned}$$

□

Fact A.2. As $\delta \rightarrow 0$,

$$\begin{aligned}\bar{\Phi}^{-1}(\delta) &= \sqrt{2 \ln 1/\delta} - \frac{\ln 2\sqrt{\pi}}{\sqrt{2 \ln 1/\delta}} \\ &\quad - \frac{\ln \ln 1/\delta}{2\sqrt{2 \ln 1/\delta}} + o\left(\frac{1}{\sqrt{\ln 1/\delta}}\right)\end{aligned}$$

Proof. The complementary error function erfc is related to $\bar{\Phi}$ via $\operatorname{erfc}(z) = 2\bar{\Phi}(\sqrt{2}z)$. Blair et al [BEJ76] show an asymptotic formula for the inverse complementary error function,

$$\begin{aligned}\operatorname{erfc}^{-1}(\delta)^2 &= \ln 1/\delta - \ln \sqrt{\pi} - \frac{1}{2} \ln(\ln 1/\delta - \ln \sqrt{\pi}) \\ &\quad + o(1) \\ \bar{\Phi}^{-1}(\delta)^2 &= 2 \ln 1/\delta - 2 \ln 2\sqrt{\pi} - \ln(\ln 1/\delta - \ln 2\sqrt{\pi}) \\ &\quad + o(1) \\ &= 2 \ln 1/\delta - 2 \ln 2\sqrt{\pi} - \ln \ln 1/\delta + o(1).\end{aligned}$$

Taking a square root,

$$\begin{aligned}&= \sqrt{2 \ln 1/\delta} \left(1 - \frac{2 \ln 2\sqrt{\pi} + \ln \ln 1/\delta + o(1)}{2 \ln 1/\delta}\right)^{1/2} \\ &= \sqrt{2 \ln 1/\delta} \left(1 - \frac{2 \ln 2\sqrt{\pi} + \ln \ln 1/\delta}{4 \ln 1/\delta} + o\left(\frac{1}{\ln 1/\delta}\right)\right) \\ &= \sqrt{2 \ln 1/\delta} - \frac{\ln 2\sqrt{\pi}}{\sqrt{2 \ln 1/\delta}} - \frac{\ln \ln 1/\delta}{2\sqrt{2 \ln 1/\delta}} \\ &\quad + o\left(\frac{1}{\sqrt{\ln 1/\delta}}\right)\end{aligned}$$

□