DISCREPANCY THEORY AND RELATED ALGORITHMS

NIKHIL BANSAL

ABSTRACT

We survey some classical results and techniques in discrepancy theory, and the recent developments in making these techniques algorithmic. The previous methods were typically based on non-constructive approaches such as the pigeonhole principle, and counting arguments involving exponentially many objects and volume of convex bodies. The recent algorithmic methods are based on an interesting interplay of methods from discrete Brownian motion, convex geometry, optimization, and matrix analysis, and their study has led to interesting new connections and progress in both discrepancy and algorithm design.

MATHEMATICS SUBJECT CLASSIFICATION 2020

Primary 68R; Secondary 05D40, 11K38, 68Q25

KEYWORDS

Discrepancy, combinatorics, algorithms, random walks

ICM INTERNATIONAL CONGRESS OF MATHEMATICIANS 2022 JULY 6—14

Proc. Int. Cong. Math. 2022, Vol. 7, pp. 5178–5210
DOI 10.4171/ICM2022/169

1. INTRODUCTION

Combinatorial discrepancy deals with the following type of question. Given a set-system (U, C) with elements $U = [n]$ and a collection $C = \{S_1, \ldots, S_m\}$ of subsets of U, how well can we color the elements red and blue so that each set $S_i \in C$ is colored as evenly as possible. Formally, let us use ± 1 to denote the colors red and blue, so that if $x(j)$ denotes the color of element j, then $|\sum_{j \in S} x(j)|$ is the imbalance for set S. Then the discrepancy of the set system (U, C) is defined as

$$\operatorname{disc}(C) = \min_{x:U \to \{-1,1\}} \max_{i \in [m]} \left| \sum_{j \in S_i} x(j) \right|,$$

that is, the minimum imbalance that must occur in at least one of the sets in C, over all possible bipartitions of U.

More generally, for an $m \times n$ matrix A, the discrepancy of A is defined as

$$\operatorname{disc}(A) = \min_{x \in \{-1,1\}^n} \|Ax\|_\infty. \tag{1.1}$$

This generalizes the definition for set systems by choosing A to be the incidence matrix for the system. Letting $v_1, \ldots, v_n$ denote the columns of A, this is the same as minimizing $\|\sum_j x(j) v_j\|_\infty$ over all ± 1 colorings x, and the problem is also referred to as vector balancing. In some settings, one also considers more general norms besides ℓ_∞, and more general objects v_i than just vectors.

Roughly speaking, discrepancy can be viewed as the study of how to divide a set of objects into two (or more) parts which are as similar as possible, with respect to various criteria. For this reason the problem arises in several applications, often in unexpected ways, and is related to various topics in mathematics and theoretical computer science [22,23,26,47,53]. For example, in computer science it has several applications in areas such as computation geometry, pseudorandomness, approximation algorithms, numerical integration, and differential privacy.

Beating random coloring. For any discrepancy problem, one option is to simply pick a random coloring by setting each $x(j)$ independently and uniformly to ± 1. However, for many problems one can do substantially better, so in a sense discrepancy theory can be viewed as the study of how to improve over the basic probabilistic method [1].

1.1. A brief history

Roughly speaking, there are three classical techniques in discrepancy. One of the earliest techniques was linear algebraic and similar to the well-known iterated-rounding technique [17,19,41]. Though this technique gives surprisingly good bounds for some problems in discrepancy, in general they are quite weak and far from optimal.

A huge breakthrough was made in the 1980s with the partial-coloring method due to Beck [18] and Spencer [63]. A similar approach based on ideas from convex geometry was developed independently by Gluskin [33]. Roughly speaking, this method guarantees the existence of a coloring of a constant fraction of the elements where every set in the set

system incurs a low discrepancy. This method is then repeated $O(\log n)$ times in order to get a full coloring of all the n elements.

The third approach was developed by Banaszczyk [4] in the late 1990s based on sophisticated ideas from convex geometry. His technique produced a full coloring directly, and led to improved bounds for many fundamental discrepancy problems.

Algorithmic aspects. Interestingly, the original proofs of the partial-coloring method and Banaszczyk's method were based on non-constructive approaches such as counting arguments, pigeonhole principle, and volume estimates of convex bodies, and did not give efficient algorithms. It was even conjectured that these results might be inherently non-algorithmic. This was problematic as in many applications of discrepancy one actually needs to be able to find good colorings efficiently.

In recent years, there has been remarkable progress in obtaining algorithmic versions of both the partial-coloring method [6, 28, 36, 45, 59] and Banaszczyk's method [8–10, 24, 35, 42]. There techniques combine ideas from linear algebra, discrete Brownian motion, optimization, and convex geometry in interesting ways, and lead to several new results and insights both in discrepancy and algorithm design. Another remarkable development has been on approximating hereditary discrepancy based on the γ_2-norm from functional analysis and semidefinite programming duality [49, 50].

In this survey, we give a brief overview of both the classical techniques and recent algorithmic results, and sketch the main ideas behind them. We also discuss some recent new directions such as online discrepancy, discrepancy of random instances, matrix discrepancy and mention several conjectures and problems that are still open. Unfortunately, we have to leave out several interesting topics, and in particular the various exciting applications of these results. We also leave out techniques for proving lower bounds for discrepancy problems.

1.2. Some examples

We describe some classical problems to give a flavor of the area, and we will use these throughout as running examples to illustrate the various techniques. Most of these problems have a long and fascinating history, that we will discuss only very briefly here.

(1) *Spencer's problem.* What is the discrepancy of an arbitrary set system with n elements and m sets?

(2) *Beck–Fiala problem.* What is the discrepancy of a set system where each element lies in at most d sets, i.e., the maximum degree is at most d?

(3) *Komlós problem.* What is the discrepancy of a matrix where the columns have ℓ_2-norm at most 1?

(4) *Prefix Komlós.* Given vectors $v_1, \ldots, v_n \in \mathbb{R}^m$ satisfying $\|v_j\|_2 \leq 1$, minimize the maximum discrepancy of prefixes, i.e., minimize $\max_{k\in[n]} \| \sum_{j=1}^{k} x(j)v_j \|_\infty$.

(5) *Tusnady's problem.* Given a set U of n arbitrary points in $[0,1]^d$, what is the discrepancy with respect to rectangles, i.e., sets $R \cap U$ where R ranges over all possible rectangles $\prod_{i=1}^{d}[a_i, b_i] \subset [0,1]^d$?

(6) *Discrepancy of Arithmetic Progressions.* Here $U = \mathbb{Z}_n$ and the sets S are of the form $S_{a,b} = \{a, a+b, a+2b, \ldots\}$. The special case of homogeneous arithmetic progressions (with $a = 0$ for each set) is called the Erdős Discrepancy Problem.

Given a coloring x and a row a, let $\mathrm{disc}(x,a) := |\sum_j x(j)a(j)|$ denote the discrepancy of x for a. For any set S, by standard probabilistic tail bounds

$$\Pr\big[\mathrm{disc}(x,S) \geq c|S|^{1/2}\big] \approx \exp\big(-c^2/2\big), \tag{1.2}$$

and thus a random coloring has discrepancy $\Omega(n^{1/2})$ or worse for all the problems above. We now describe the various improved bounds known for them. We shall give the details in later sections.

1.2.1. Spencer's problem

For an arbitrary set system, (1.2) and a union bound over the m sets implies that a random coloring has discrepancy $O((n \log m)^{1/2})$ with high probability (whp). In an influential work, Spencer [63] and Gluskin [33] showed the following result.

Theorem 1.1. *Any set system with $m \geq n$ sets has discrepancy $O((n \log 2(m/n))^{1/2})$. For $m \leq n$, the discrepancy is $O(m^{1/2})$.*

In particular, for $m = n$ this gives $O(n^{1/2})$ discrepancy, which is also the best possible, answering a question of Erdős. While this $O(\log n)^{1/2}$ factor improvement over the random coloring may seem relative minor, Spencer developed the partial coloring method to prove Theorem 1.1, which has become a key tool and gives huge improvements for many other problems.

1.2.2. Beck–Fiala and Komlós problem

Beck and Fiala [19], in one of the first applications of the iterated rounding technique, showed the following result.

Theorem 1.2. *Any set system with maximum degree d has discrepancy at most $2d-1$.*

A long-standing conjecture is the following.

Conjecture 1.2.1 ([19]). *The discrepancy of any set system with degree d is $O(d^{1/2})$.*

If we allow a mild dependence on n, the partial-coloring method gives a bound of $O(d^{1/2}\log n)$. The best known bound in this direction is $O((d\log n)^{1/2})$ due to Banaszczyk [4], based on a more general result that we shall see later.

Scaling the entries by $d^{-1/2}$, notice that the Beck–Fiala problem is a special case of the Komlós problem. For the Komlós problem, the partial-coloring method gives an $O(\log n)$

bound, and Banaszczyk's method gives the best known bound of $O((\log n)^{1/2})$. The following conjecture generalizes Conjecture 1.2.1.

Conjecture 1.2.2 (Komlós). *Given any $v_1, \ldots, v_n \in \mathbb{R}^m$ satisfying $\|v_j\|_2 \leq 1$ for all j, there is an $x \in \{-1, 1\}^n$ such that $\| \sum_j x(j) v_j \|_\infty = O(1)$.*

1.2.3. Prefix discrepancy

The study of discrepancy problems involving prefixes of a given sequence of vectors also has a long history and several surprising connections to other classical ordering problems. See, e.g., [17] for a fascinating survey and also [5, 20, 27, 47]. We restrict our focus here to the prefix version of the Komlós problem. The best bound known here is $O((\log n)^{1/2})$ due to Banaszczyk [5], where he further extended his method from [4] to handle prefixes.

Given this extension, a natural question is whether the prefix Komlós problem is any harder than the one without prefixes.

Problem 1.3. Is the discrepancy of the prefix Komlós problem $O(1)$?

There is no clear consensus here, and in fact for some discrepancy problems it is known that considering prefixes makes the problem harder [30, 52].

Algorithmic aspect. As we shall see later, there are several algorithmic approaches known by now for the partial-coloring method and for Banaszczyk's method in [4] without prefixes. However, the best algorithmic bound we know for the prefix version is still $O(\log n)$, based on partial coloring approach, and the following question is very interesting.

Problem 1.4. Find an efficient algorithm to obtain an $O((\log n)^{1/2})$ discrepancy coloring for the prefix Komlós problem.

1.2.4. Tusnady's problem

Here one can do exponentially better than random colorings, and these ideas have significant applications in numerical integral and quasi-Monte Carlo methods [23, 47].

The case of $d = 2$ is already instructive to see the relative power of various techniques. Moreover, we still do not know the right answer here. Linear algebraic methods give a bound of $O(\log^4 n)$. Using partial coloring, this can be pushed to about $O(\log^{5/2} n)$ [47]. The current best bound is $O(\log^{3/2} n)$ due to Nikolov [54], based on Banaszczyk's result for prefix discrepancy [5]. On the other hand, the best known lower bound is $\Omega(\log n)$ [49, 61].

For general d, after a long line of work, the current lower and upper bounds are $\Omega_d(\log^{d-1} n)$ [49] and $O_d(\log^{d-1/2} n)$ [54]. Closing this gap is an important open problem.

Conjecture 1.4.1. *The discrepancy of Tusnady's problem in d dimensions is $O(\log^{d-1} n)$.*

1.2.5. Arithmetic progressions

This problem has a long history, including results of Weyl [69] and Roth [58]. Using Fourier analysis, Roth [58] proved a lower bound of $\Omega(n^{1/4})$. Interestingly, Roth believed that

his result might not be best possible and the right exponent might be $1/2$, as suggested by random colorings. Eventually, Matoušek and Spencer [51] gave a matching $O(n^{1/4})$ upper bound using the partial coloring method.

For homogeneous arithmetic progressions, an $O(\log n)$ upper bound follows from a simple explicit construction. A famous question of Erdős was whether the discrepancy is $O(1)$. This was answered negatively in a breakthrough work by Tao [67].

1.3. Hereditary discrepancy and rounding

An important application of discrepancy is in rounding a fractional solution to an integral one without introducing much error, based on the following result of Lovász, Spencer, and Vesztergombi [44].

Theorem 1.5 ([44]). *For any $x \in \mathbb{R}^n$ satisfying $Ax = b$, there is a $\tilde{x} \in \mathbb{Z}^n$ with $\|\tilde{x} - x\|_\infty < 1$, such that $\|A(x - \tilde{x})\|_\infty \leq \operatorname{herdisc}(A)$.*

Here $\operatorname{herdisc}(A)$ is the *hereditary discrepancy* of A, which is a more robust version of discrepancy, and defined as the maximum discrepancy over all column restrictions of A,

$$\operatorname{herdisc} A = \max_{S \subset [n]} \operatorname{disc}(A_{|S}) = \max_{S \subset [n]} \min_{x \in \{-1,1\}^n} \|Ax\|_\infty.$$

For most classes of set systems, any upper bound on discrepancy is also a bound on hereditary discrepancy, as the class itself may be closed under taking subset of columns. For example, this holds for all the problems in Section 1.2, except for the case of arithmetic progressions, which is an example of a particular set system.

Rounding via discrepancy. To see the idea behind Theorem 1.5, suppose that x is $1/2$-integral (i.e., each $x(j)$ has fractional part 0 or $1/2$). Let S be the set of variables with fractional part $1/2$, and let y be ± 1 coloring of S with discrepancy $\operatorname{disc}(A_{|S})$. Then $x' = x + y/2$ is integral and

$$\|Ax' - Ax\|_\infty = \|A(y/2)\|_\infty = \operatorname{disc}(A_{|S})/2 \leq \operatorname{herdisc}(A)/2.$$

That is, the signs of y are used to decide whether to round each $x(j)$ up or down. For arbitrary x, Theorem 1.5 follows by applying this to each bit after the decimal starting from the least significant bit.

The problem of rounding arises naturally for example in designing efficient approximation algorithms for discrete optimization problems. However, note that Theorem 1.5 only shows the existence of a good rounding, and gives no clue on how to actually find one efficiently. We shall see an algorithmic version of Theorem 1.5 in Section 3.1. In general, the recent algorithmic progress on discrepancy has led to several new results in approximation algorithms, a particularly notable result is [60].

2. CLASSICAL TECHNIQUES

We now describe the classical techniques of (i) the linear algebraic method, (ii) partial coloring, and (iii) Banaszczyk's method. Interestingly, the linear algebraic idea will also

play a key role in many of the algorithmic versions of partial coloring and Banaszczyk's method that we shall see later in Sections 3 and 4.

2.1. Linear algebraic method

This technique is simple but it can be surprisingly powerful. It is widely used in combinatorial optimization and is also referred to as iterated rounding [41].

For discrepancy problems it works as follows. Let $B \in \mathbb{R}^{m \times n}$ be the input matrix. The algorithm starts with the all-zero coloring $x_0 = (0, 0, \ldots, 0)$, and updates the coloring over several iterations $t = 1, 2, \ldots, T$, until the final coloring $x_T \in \{-1, 1\}^n$. The intermediate colorings satisfy $x_t \in [-1, 1]^n$, and once some color reaches ± 1 it is never updated again.

It remains to specify how the coloring is updated in each iteration. Call a variable j *alive* at time t, if $x_{t-1}(j) \in (-1, 1)$ and let A_t be the set of alive variables at the beginning of time t. The idea is to pick a suitable subset B_t of rows of B, with $\text{rank}(B_t) < |A_t|$, and consider some nonzero solution v_t satisfying (i) $B_t v_t = 0$ and (ii) $v_t(j) = 0$ for $j \in [n] \setminus A_t$. Such a solution exists as there are $|A_t|$ alive variables, and $\text{rank}(B_t) \leq |A_t|$.

The coloring is updated as $x_t = x_{t-1} + \delta v_t$, where $\delta > 0$ is chosen so that x_t stays in $[-1, 1]^n$ and at least one more color reaches ± 1 compared to x_{t-1}. The ingenuity lies in choosing B_t at each time t.

Let us see how to use this template to prove the Beck–Fiala theorem.

Theorem 1.2. *Any set system with maximum degree d has discrepancy at most $2d - 1$.*

Proof. Let B denote the incidence matrix of the set system. By our assumption, each column of B has at most d ones. Let us apply iterated rounding, where at iteration t we choose B_t to consist of rows S_i with $|A_t \cap S_i| > d$. Call such rows *large*. As each column of B has at most d ones, the number of ones in B restricted to columns in A_t is at most $d|A_t|$, and so the number of large rows is strictly less than $|A_t|$ and thus $\text{rank}(B_t) < |A_t|$.

To bound the final discrepancy, notice that as long as a row is large, its discrepancy stays 0. But once it has at most d alive elements, then no matter how these variables get updated in subsequent iterations, the additional discrepancy must be strictly less than $2d$ (e.g., if all the d alive variables were all -0.999 but get set to 1 eventually). As the final discrepancy of a set system is integral, this gives the bound of $2d - 1$. ∎

For more ingenious applications of this method to discrepancy, we refer to the survey by [17] and references therein.

2.2. Partial coloring method

We now describe the partial coloring lemma of Spencer and give some applications. We then describe the convex geometric proof of this result due to Gluskin [33] based on an exposition of Giannopoulos [32].

Theorem 2.1 (Partial coloring lemma). *Let A be an $m \times n$ matrix with rows $a_1, \ldots, a_m$. For each $i \in [m]$, let $\Delta_i = \lambda_i \|a_i\|_2$ be target discrepancy bound for row i. If the λ_i satisfy*

$$\sum_{i \in [m]} g(\lambda_i) \leq n/5, \tag{2.1}$$

where

$$g(\lambda) = \begin{cases} K \exp(-\lambda^2/9) & \text{if } \lambda > 0.1, \\ K \ln(\lambda^{-1}) & \text{if } \lambda \leq 0.1, \end{cases}$$

and K is some absolute constant. Then there exists $x \in \{-1, 0, 1\}^n$ with $|\{j : |x(j)| = 1\}| \geq n/10$ and $\operatorname{disc}(y, a_i) \leq \Delta_i$ for each $i \in [m]$.

Comparison with union bound. It is instructive to compare this with the standard union bound argument. For a random coloring x, as $\Pr[\operatorname{disc}(a_i, x) \geq \lambda \|a_i\|_2] \approx \exp(-\lambda_i^2/2) \approx g(\lambda_i)$. For the union bound to work, we need to choose λ_i to (roughly) satisfy the condition $\sum_i g(\lambda_i) < 1$. In contrast, Lemma 2.1 allows $\sum_i g(\lambda_i) = \Omega(n)$. This gives substantially more power. For example, suppose A is a 0–1 matrix corresponding to a set system. The union bound argument cannot ensure that $\Delta_i \ll |S_i|^{1/2}$ for even a couple of sets, while Theorem 2.1 allows us to set $\Delta_i < 1$ for $O(n/\log n)$ sets. As $x \in \{-1, 0, 1\}^n$, this in fact gives a partial coloring with exactly zero discrepancy for those sets!

2.2.1. Applications

The partial coloring method is very general and is widely used in discrepancy theory. We show how it directly gives Theorem 1.1 and the $O(d^{1/2} \log n)$ bound for the Beck–Fiala problem.

Proof of Theorem 1.1 for Spencer's problem. Let us assume that $m \geq n$. The case of $m \leq n$ follows by reducing $n = m$ by a standard linear algebraic trick. The coloring is constructed in phases. Let $n_0 = n$ and let n_k be the number of uncolored elements in phase k. In phase k, we apply Theorem 2.1 to the set system restricted to these n_k elements with $\Delta_k = c(n_k \log(2m/n_k))^{1/2}$ for each row, and verify that (2.1) holds for large enough $c = O(1)$. This gives a partial coloring on $\geq n_k/10$ elements, so $n_k \leq (0.9)^k n$ and summing up over the phases, total discrepancy is at most $\Delta_0 + \Delta_1 + \cdots = O((n \log(m/n))^{1/2})$.

***O*($d^{1/2} \log n$) discrepancy for the Beck–Fiala problem.** Again, the coloring is constructed in phases where $n_k \leq n(0.9)^k$ elements are uncolored in phase k. In phase k, let $s_{k,j}$ denote the number of sets with number of uncolored elements in the range $[2^j, 2^{j+1})$. Then $s_{k,j} \leq \min(m, n_k d/2^j)$ as the degree is at most d. A simple computation shows that (2.1) holds for $\Delta_i = cd^{1/2}$ for each i, for large enough $c = O(1)$. The result then follows directly as there are $O(\log n)$ phases and each set incurs $O(d^{1/2})$ discrepancy in each phase.

2.2.2. Proof of the partial coloring lemma

The original proof of Spencer was based on the pigeonhole principle and the entropy method and has several nice expositions, e.g., [1]. We sketch here the convex geometric proof of Gluskin [33].

The simple observation that ties discrepancy to geometry is the following.

Observation 2.1.1. *For a $m \times n$ matrix A, there is a coloring with discrepancy at most Δ_i for row a_i iff the polytope $P = \{x : |a_i x| \leq \Delta_i, i \in [m]\}$, contains some point in $\{-1, 1\}^n$.*

Similarly, Theorem 2.1 is equivalent to showing that the polytope P contains some point in $\{-1, 0, 1\}^n$ with at least $n/10$ nonzero coordinates, if the Δ_i satisfy (2.1).

Gluskin relates this property to the Gaussian volume of P. Call a convex body K symmetric if $x \in K$ implies $-x \in K$. Let $\gamma_n(x) = (2\pi)^{-n/2} \exp(-\|x\|_2^2/2)$ denote the standard n-dimensional Gaussian measure. For ease of exposition, we ignore the constants in Theorem 2.1.

Theorem 2.2 (Gluskin). *There is a small constant $\delta > 0$, such that any symmetric convex body $K \in \mathbb{R}^n$ with $\gamma_n(K) \geq 2^{-\delta n}$ contains $y \in \{-1, 0, 1\}^n$ with at least δn coordinates ± 1.*

The proof is based on a nice counting argument.

Proof. For $x \in \mathbb{R}^n$, let $K_x := K + x$ denote K shifted by x. As K is symmetric, we have that $\gamma_n(K_x) \geq \exp(-\|x\|^2/2)\gamma_n(K)$, as the densities of any two symmetric points y and $-y$ upon shifting by x satisfy,

$$\gamma_n(y + x) + \gamma_n(x - y) \geq 2\big(\gamma_n(x - y)\gamma_n(y + x)\big)^{1/2} = 2\exp\big(-\|x\|^2/2\big)\gamma_n(y).$$

Consider the 2^n copies K_x for all $x \in \{-1, 1\}^n$. As the total Gaussian volume of these copies is at least $2^n \exp(-n/2)\gamma_n(K) = 2^{cn}$, for some $c > 0$, there exists some point z contained in at least 2^{cn} copies. So there must exist some $x, x' \in \{-1, 1\}^n$ differing in $\Omega(n)$ coordinates such that z lies in both K_x and $K_{x'}$. Suppose $z = k_1 + x = k_2 + x'$ for some $k_1, k_2 \in K$. Then $y := (x - x')/2 = (k_2 - k_1)/2 \in K$ as K is symmetric and convex, and, moreover, $y \in \{-1, 0, 1\}^n$ with $\Omega(n)$ coordinates ± 1. ∎

Theorem 2.1 follows by relating the condition (2.1) on λ_i to the volume $\gamma_n(P)$.

Gaussian measure of polytopes. For a vector $a \in \mathbb{R}^n$ and scalar $b > 0$, define the slab $S(a, b) = \{x : |\langle x, a\rangle| \leq b\}$. Then $S(a, \lambda\|a\|_2)$ is symmetric and convex, with measure $\gamma_n(S(a, b)) = \gamma_1([-\lambda, \lambda])$, and $P = \bigcap_{i=1}^m S(a_i, \lambda_i\|a_i\|_2)$ is an intersection of slabs.

By the Sidak–Khatri lemma (see, e.g., [33]), for any symmetric convex body K and slab S, we have that $\gamma_n(K \cap S) \geq \gamma_n(K)\gamma_n(S)$, and hence $\gamma_n(P) \geq \prod_i \gamma_n(S(a_i, \lambda_i\|a_i\|_2)$. As $\gamma_1([-\lambda, \lambda]) \approx 1 - O(\exp(-\lambda^2/2))$ for $\lambda \geq 1$ and $O(\lambda)$ for $\lambda < 1$, we have that $\log(\gamma_1[-\lambda, \lambda]) \approx -g(\lambda)$, and thus condition (2.1) implies that $\gamma_n(P) \geq 2^{-\delta n}$.

2.3. Banaszczyk method

A problem with the partial coloring method is that it requires $O(\log n)$ rounds to obtain a full coloring, which can result in an extra $O(\log n)$ factor loss in the discrepancy bound, as we saw for the Beck–Fiala problem. The following result of Banaszczyk [4] gives a way to find a full coloring directly and can give better results. The form of the result also makes it broadly applicable in other settings.

Theorem 2.3 ([4]). *Given any convex body $K \subseteq \mathbb{R}^m$ of Gaussian measure $\gamma_m(K) \geq 1/2$, and vectors $v_1, \ldots, v_n \in \mathbb{R}^m$ of ℓ_2 norm at most $1/5$, there exists a coloring $x : [n] \to \{-1, 1\}$ such that $\sum_{j=1}^n x(j)v_j \in K$.*

While this statement looks similar to Theorem 2.1, a crucial difference is that the Gaussian measure and convex body K here are in the *output space* and are m-dimensional, while K in Theorem 2.1 is in the *input space* and n-dimensional.

The proof of Theorem 2.3 involves some delicate computation and a non-trivial idea of Ehrhard symmetrization. However, the main idea is very clean that we sketch below.

Proof. The key step is to show that for any convex body K with $\gamma_m(K) \geq 1/2$ and $u \in \mathbb{R}^m$ with $\|u\|_2 \leq 1/5$, there is a convex body $K * u$ contained in $(K - u) \cup (K + u)$ such that $\gamma_m(K * u) \geq 1/2$.

Given this fact, Theorem 2.3 follows by induction on the number of vectors n. It trivially holds for $n = 0$ as $\mathbf{0} \in K$ as $\gamma_m(K) \geq 1/2$. Suppose inductively that it holds for some $n - 1$. Consider the convex body $K' = K * v_n$. As $\gamma_m(K') \geq \gamma_m(K) \geq 1/2$, by induction there exist $x(1), \ldots, x(n-1) \in \{-1, 1\}$ such that $u = x(1)v_1 + \cdots + x(n-1)v_{n-1} \in K'$. But as $K' \subset (K - v_n) \cup (K \cup v_n)$, at least one of $u + v_n$ or $u - v_n$ must lie in K, giving the sign $x(n)$ such that $u + x(n)v_n = \sum_{i=1}^n x(i)v_i \in K$. ∎

Bound for the Komlós problem. Theorem 2.3 directly gives the $O((\log n)^{1/2})$ bound for the Komlós problem. This follows as $\mathbb{E}\|g\|_\infty = O((\log m)^{1/2})$ for a random Gaussian vector g in $\mathbb{R}^m$, and so choosing K to be $2\mathbb{E}\|g\|_\infty$ times the unit ℓ_∞-ball in $\mathbb{R}^m$, by Markov's inequality we have that $\gamma_m(K) \geq 1/2$. As Theorem 2.3 requires $\|v\|_2 \leq 1/5$, we can further scale K by a factor of 5. Finally, we can assume $m \leq n^2$, as $\|v_j\|_2 \leq 1$ for each j implies that at most n^2 rows a_i can have $\|a_i\|_1 \geq 1$.

Banaszczyk's theorem for prefix discrepancy. In a subsequent work, Banaszczyk [5] further extended this result to handle prefixes, using a clever inductive argument.

Theorem 2.4 ([5]). *Given vectors $v_1, \ldots, v_n \in \mathbb{R}^m$ of ℓ_2 norm at most $1/5$ and any convex body $K \subseteq \mathbb{R}^m$ with $\gamma_m(K) \geq 1 - 1/(2n)$, there exists a coloring $x : [n] \to \{-1, 1\}$ such that each for $k = 1, \ldots, n$, the prefix sum satisfies $\sum_{j=1}^k x(j)v_j \in K$.*

Proof. Consider the sequence of symmetric convex bodies K_j defined iteratively as $K_n = K$ and $K_j = (K_{j+1} * v_{j+1}) \cap K$, for $j = n-1, \ldots, 1$. We first show that $\gamma_m(K_j) \geq 1 - (n - j + 1)/2n$ for $j \in [n]$ by backwards induction. Indeed, $\gamma_m(K_n) = \gamma_m(K) \geq 1 - 1/(2n)$ in the base case. If this it holds for some $j \leq n$, then

$$\begin{aligned}\gamma_m(K_{j-1}) &= \gamma_m(K_j * v_j) \cap \gamma_m(K) \geq \gamma_m(K_j * v_j) - \big(1 - \gamma_m(K)\big) \\ &\geq \gamma_m(K_j) - \big(1 - \gamma_m(K)\big) \geq 1 - \frac{n - j + 1}{2n} - \frac{1}{2n} = 1 - \frac{n - j}{2n},\end{aligned}$$

where we use that $\gamma_m(K_j * v_j) \geq \gamma_m(K_j)$ as $\gamma_m(K_j) \geq 1/2$.

So $\gamma_m(K_1) \geq 1/2$ and K_1 is convex, and a simple calculation shows that either v_1 or $-v_1$ lies in K_1. We now apply induction in the forward direction. Suppose there is some

$j \geq 1$ such that there are signs $x(1), \ldots, x(j)$ satisfying (i) $u := \sum_{i=1}^{j} x(i)v_i \in K_j$ and (ii) $\sum_{i=1}^{k} x(i)v_i \in K$ for all $k \leq j$. To continue the induction, we need to show that (i) $u \in K$ and (ii) that there is a sign $x(j+1)$ such that $u + x(j+1)v_{j+1} \in K_{j+1}$. Now, $u \in K$ clearly holds as by (i) we have $u \in K_j \subset K$. Now, for the sake of contradiction suppose that both $u + v_{j+1}$ and $u - v_{j+1} \notin K_{j+1}$. Then $u \notin K_{j+1} + v_{j+1} \cup K_{j+1} - v_{j+1}$ and hence $u \notin K_{j+1} * v_{j+1}$. By definition, as $K_j = K \cap (K_{j+1} * v_{j+1}) \subset K_{j+1} * v_{j+1}$, this contradicts our inductive assumption that $u \in K_j$. ■

Bound for prefix Komlós. Theorem 2.4 directly implies an $O((\log n)^{1/2})$ discrepancy for the prefix version of the Komlós problem. In particular, the condition that $\gamma_m(K) \geq 1 - 1/(2n)$ instead of $\geq 1/2$ in Theorem 2.3 make no difference beyond a constant factor as $\Pr[\|g\|_\infty \geq \mathbb{E}\|g\|_\infty + t] \leq \exp(-t^2/2)$ by concentration for Lipschitz functions of Gaussians, and choosing $t = O((\log n)^{1/2})$.

3. ALGORITHMS FOR PARTIAL COLORING

In the next few sections we describe the progress on making these results algorithmic. We first describe several different algorithmic proofs for partial coloring. In Section 4 we describe the algorithmic approaches for Banaszczyk's method as stated in Theorem 2.3. In Section 5, we describe an algorithm to approximate the hereditary discrepancy of any arbitrary matrix.

The algorithms for partial coloring can be divided into two types: either based on a random walk approach, or a direct optimization based approach.

Random-walk based approaches. Bansal [6] gave the first algorithm for various applications of partial coloring such as the $O(n^{1/2})$ bound for Spencer's problem with $m = O(n)$ sets and the $O(d^{1/2} \log n)$ bound for the Beck–Fiala problem. Subsequently, Lovett and Meka [45] designed an elegant and substantially simpler algorithm that gave an algorithmic version of the full partial coloring lemma as stated in Theorem 2.1.

These algorithms can be viewed as a randomized version of the iterated rounding method, where one starts with the all zero-coloring, and updates the variables gradually using a correlated Brownian motion with small discrete steps. The variables are fixed once they reach ± 1, and correlations between the variables are chosen to ensure that each row has low discrepancy. Bansal's algorithm was based on solving a suitable semidefinite program (SDP) at each time step to generate the covariance matrix for the random walk. Lovett and Meka showed that one can simply do a standard discrete Brownian motion in the subspace orthgonal to tight discrepancy constraints, without the need to solve any SDPs.

Direct methods. Later, Rothvoss [59] further extended the result of Lovett and Meka from polytopes to general symmetric convex bodies and gave an algorithmic version of Theorem 2.2. His algorithm is extremely elegant and simple to describe. A related algorithm was given by Eldan and Singh [28]. Both these algorithms are based on solving a very simple optimization problem.

We now describe these algorithms and sketch the main ideas behind their analysis.

3.1. The SDP-based approach

We start with the SDP-based approach. Even though the latter algorithms are more general and simpler, this approach is very natural and motivates why the Brownian motion is needed. It is also the only approach we know for some problems such as the algorithmic version of Theorem 1.5, that we describe below in Theorem 3.1. More importantly, it is quite flexible and can be extended in various ways by adding new SDP constraints, as we shall see later in Section 4.

A relaxation for discrepancy. Given an input matrix A, a natural approach to find a low discrepancy coloring for it is to first solve some convex programming relaxation and then try to round the solution suitably to ± 1. Let us first consider linear programming relaxations.

Recall that a linear program (LP) consists of variables $x_1, \ldots, x_n \in \mathbb{R}$, and the goal is to optimize some linear objective $c^T x$ subject to linear constraints $a_i^T x \leq b_i$ for $i \in [m]$. LPs can be solved optimally in time polynomial in n, m, and the bit length of the input.

Let a_i denote the ith row of A, then the natural LP relaxation for discrepancy is,

$$\min t \quad \text{s.t.} \quad -t \leq a_i x \leq t, \ \ \forall i \in [m] \ \text{ and } \ -1 \leq x_j \leq 1, \ \ \forall j \in [n].$$

However, this always has the trivial solution $x = \mathbf{0}$ with objective $t = 0$, which is useless.

So let us consider a more general class of optimization problems called semidefinite programs (SDPs). An SDP can be viewed as an LP with variables of the form x_{ij} for $1 \leq i, j \leq n$, arranged as entries of an $n \times n$ matrix X, where we require that X be symmetric and positive semidefinite, denoted by $X \succeq 0$. For matrices A, B, let $\langle A, B \rangle = \mathrm{Tr}(A^T B) = \sum_{ij} A_{ij} B_{ij}$ denote the trace inner product. An SDP is an optimization problem of the form

$$\max \langle C, X \rangle \quad \text{s.t. } \langle A_k, X \rangle \leq b_k, \ 1 \leq k \leq m, \ X \succeq 0,$$

where $C, A_1, \ldots, A_m \in \mathbb{R}^{n \times n}$.

SDPs can be solved to any desired level of accuracy in polynomial time. As $X \succeq 0$ iff it is the Gram matrix of some vectors $w_1, \ldots, w_n \in \mathbb{R}^n$, i.e., $X_{ij} = \langle w_i, w_j \rangle$, SDPs can be viewed as *vector programs* where the variables are the vectors w_i and we can impose any linear constraints on their inner products (but not on the w_i themselves).

SDP relaxation for discrepancy. Let λ be some upper bound on the discrepancy $\mathrm{disc}(A)$, and consider the following SDP:

$$\left\| \sum_j a_{ij} w_j \right\|_2^2 \leq \lambda^2 \quad \text{for } i \in [m], \ \|w_j\|_2^2 = 1, \ i \in [n],$$

Let us call a feasible solution to this SDP a vector-coloring for A, and the smallest λ for which it is feasible as the vector discrepancy, $\mathrm{vecdisc}(A)$. Clearly, $\mathrm{vecdisc}(A) \leq \mathrm{disc}(A)$.

At first glance, this SDP also does not seem useful. For example, for Spencer's problem, the solution $w_i = e_i$, where w_i is the ith standard basis vector, is always feasible with $\lambda = n^{1/2}$, irrespective of the matrix A. However, this SDP becomes quite useful

when $\lambda \ll n^{1/2}$ as it gives nontrivial correlations between the vectors w_i that we can exploit. Below, we describe a very simple algorithm that gives the following algorithmic version of Theorem 1.5.

Theorem 3.1 ([6]). *Given any $A \in \mathbb{R}^{m \times n}$, there is an efficient algorithm that, with high probability, finds a coloring with discrepancy $O((\log m \log n)^{1/2}\, \mathrm{herdisc}(A))$.*

Later, we will also describe a simple $O((n \log\log\log n)^{1/2})$ bound for Spencer's problem with $m = O(n)$ using this SDP. This is perhaps surprising, as apriori the naive solution $w_i = e_i$ does not give any meaningful correlations between the elements and corresponds to random coloring.

3.1.1. Algorithm for Theorem 3.1

Before describing the algorithm, it is instructive to see why a direct approach for rounding the SDP does not work.

Problem with direct rounding. For simplicity, let us suppose that $\lambda = 0$ for some matrix A. Then the vectors $w_1, \ldots, w_n$ produced by the SDP solution are nicely correlated so that $\sum_j a_{ij} w_j = 0$ for each row i.

To convert the w_j into scalars while preserving the correlations, let us pick a random Gaussian vector $g \in \mathbb{R}^n$, with each coordinate $g_k \sim N(0,1)$ independently and project the vectors w_j on g to obtain $y_j = \langle w_j, g\rangle$. Then as the g_k are iid $N(0,1)$, we have that $\langle g, w\rangle = \sum_k g_k w(k) \sim N(0, \|w\|_2^2)$ for any vector $w \in \mathbb{R}^n$, and hence (i) $y_j \sim N(0,1)$ for each j as $\|w_j\|_2^2 = 1$ and (ii) $\sum_j a_{ij} y_j = 0$ for each row i. This seems very close to what we want except that $y_j \sim N(0,1)$ instead of ± 1.

However, the following hardness result of Charikar, Newman, and Nikolov [21] rules out any reasonable way of rounding these y_j to ± 1.

Theorem 3.2 ([21]). *Given a set system on n elements and $m = O(n)$ sets, it is NP-hard to distinguish whether it has discrepancy 0 or $\Omega(\sqrt{n})$.*

In particular, this implies that there must exist set systems with discrepancy $\Omega(\sqrt{n})$ but vector-discrepancy 0 (otherwise solving the SDP would give an efficient way to distinguish between set systems with discrepancy 0 and $\Omega(\sqrt{n})$).

Discrete Brownian motion. So instead of trying to round the y_j's directly to ± 1, the algorithm will gradually obtain a ± 1 coloring by combining solutions of various SDPs over time. We first give a overview of the algorithm.

More precisely, at time 0, we start with the coloring $x_0 = (0, \ldots, 0)$ and modify it over time as follows. Let x_{t-1} denote the fractional coloring at time $t-1$. Then $x_t = x_{t-1} + \Delta x_t$ is obtained by adding a small update vector Δx_t to x_{t-1}. As the perturbations are added, the colors evolve over time, and once a color reaches ± 1 it is frozen and no longer updated. The updates Δx_t are obtained by solving the SDP with $\lambda = \mathrm{herdisc}(A)$, restricted to the alive elements and setting $\Delta x_t(j) = \gamma \langle g, w_j\rangle$, where g is a random gaussian and γ is a small multiplier.

Formal description. Let $\gamma = \max_{i,j} |A_{ij}|/n$ and $\ell = 8 \log n/\gamma^2$. Let x_t and A_t denote the coloring and the set of alive (unfrozen) variables at the end of time t. Let $\lambda = \text{herdisc}(A)$.

(1) Initialize $x_0(j) = 0$ for $j \in [n]$ and $A_0 = \emptyset$.

(2) At each time step $t = 1, 2, \ldots, \ell$, do the following.
Solve the following SDP:

$$\left\| \sum_j a_{ij} w_j^t \right\|_2^2 \leq \lambda^2 \quad \forall i \in [m], \quad \|w_j^t\|_2^2 = 1 \text{ if } j \in A_{t-1}, \text{ else } \|w_j^t\|_2^2 = 0.$$

Pick a random Gaussian $g_t \in \mathbb{R}^n$, and set $x_t(j) = x_{t-1}(j) + \gamma \langle g_t, w_j^t \rangle$.
Set $A_t = \{j : |x_t(j)| < 1\}$.

(3) Set $x_\ell(j) = -1$ if $x_\ell(j) < -1$ and $x_\ell(j) = 1$ otherwise. Output x_ℓ.

Analysis. We now sketch the ideas behind Theorem 3.1. First, notice that as $\lambda = \text{herdisc}(A)$, the SDP above is always feasible no matter which variables are alive.

Let us now see how the colors of the elements and the discrepancies of the rows evolve over time. Fix some element j. Its color $x_t(j)$ starts at 0 at $t = 0$ and evolves as a martingale with updates $\Delta x_t(j) = \gamma \langle w_j^t, g_t \rangle$ until it is frozen. As $\|w_j^t\| = 1$, we have $\Delta x_t(j) \sim N(0, \gamma^2)$ and thus $x_t(j)$ will reach ± 1 in $O(1/\gamma^2)$ steps with constant probability. As there are $\ell = O(\log n/\gamma^2)$ steps, whp all elements will reach ± 1, by the end of the algorithm.

Now fix some row i. Its discrepancy $x_t(a_i) := \sum_j a_{ij} x_t(j)$ is 0 at $t = 0$, and evolves as $\sum_j a_{ij} \Delta x_t(j) = \sum_j \gamma \langle g_t, \sum_j a_{ij} w_j^t \rangle$ at step t. As $\| \sum_j a_{ij} w_j^t \|^2 \leq \lambda^2$, the sequence $x_t(a_i)$ forms a martingale with Gaussian increments with variance at most $\gamma^2 \lambda^2$. As $\ell = O(\log n/\gamma^2)$, by standard martingale concentration and union bound over the m constraints, each row has final discrepancy $O(\ell^{1/2} \cdot \gamma\lambda \cdot (\log m)^{1/2}) = O(\lambda (\log m \log n)^{1/2})$ whp.

Finally, whp truncating $x_\ell(j)$ to ± 1 introduces negligible error for any row. This follows as $\Delta x_t(j) \sim N(0, \gamma^2)$ we have that whp $|x_t(j)| < 1 + \gamma \cdot O((\log n)^{1/2})$ when it freezes. As $\text{herdisc}(A) \geq \max_{ij} |A_{ij}|$ and $\gamma = \max_{ij} |A_{ij}|/n \leq \text{herdisc}(A)/n$, the rounding error is negligible.

3.1.2. Algorithmic version of Spencer's result

The above approach is quite flexible, e.g., the discrepancy bounds λ_i^t for each row i and be chosen adaptively at time t. We describe a simple version of this idea that already gives a $\beta n^{1/2}$ for $\beta = c(\log \log \log n)^{1/2}$ bound for Spencer's problem with $m = n$ sets, and thus beats random coloring.

As in Section 2.2.1, it suffices to obtain a partial coloring with $O(\beta n^{1/2})$ discrepancy. Let us run the algorithmic template above for $\ell = 100/\gamma^2$ steps, using the following

SDP relaxation for partial coloring at each time t:

$$\left\| \sum_{j \in S_i} w_j \right\|_2^2 \leq \lambda_i^2 \quad \text{for } i \in [m], \tag{3.1}$$

$$\sum_{j \in A_{t-1}} \|w_j\|_2^2 \geq |A_{t-1}|/10, \tag{3.2}$$

$$\|w_j\|_2^2 \leq 1 \quad \forall j \in A_{t-1}, \text{ else } \|w_j\|^2 = 0.$$

Notice that the λ_i on the right hand side in (3.1) can be different for each rows. The constraint (3.2) says that at least $|A_{t-1}|/10$ elements must be colored.

The bounds λ_i are set as follows. Initially, $\lambda_i = cn^{1/2}$ for each S_i where c is a large enough constant. If the discrepancy $|x_t(S_i)|$ for S_i exceeds $\beta n^{1/2}$ at any time, we label S_i *dangerous* and set $\lambda_i^2 = n/\log n$ at all future time steps.

The result follows from the following two observations.

Lemma 3.3. *If the SDPs are feasible at all time steps, then whp each set has discrepancy $O(\beta n^{1/2})$, and at least $\Omega(n)$ elements are colored ± 1 at the end of the algorithm.*

Proof. (Sketch) By the choice of the λ_i, once a set becomes dangerous, its discrepancy evolves as a martingale with Gaussian increments with variance at most $\gamma^2 n/\log n$. As there at most $\ell = O(\gamma^{-2})$ time steps, whp each set incurs an additional discrepancy of at most $O(n^{1/2})$.

Next, the variance $\mathbb{E}[\Delta x_t(j)^2]$ increases by at least $\gamma^2/10$ on average for the alive variables at each step t by the constraint (3.2). As $\ell = 100\gamma^{-2}$, a simple Markov argument shows that a constant fraction of the elements will reach ± 1 with at least constant probability. ∎

Lemma 3.4. *With probability $1 - o(1)$, all the SDPs are feasible.*

Proof. (Sketch) As $\lambda_i \leq O(n^{1/2})$ at each time and $\ell = O(\gamma^{-2})$, each set S_i has discrepancy $O(n^{1/2})$ in expectation. So by standard martingale concentration, with probability $1 - o(1)$, the fraction of sets that ever become dangerous $2\exp(-\Omega(\beta^2)) \ll (\log\log n)^{-2}$ for c large enough. Let us condition on this event. We will show that the SDP is feasible at each step using Theorem 2.1. Indeed, as each dangerous set S_i contributes $g(\Delta_i/|S_i|^{1/2}) \leq g(1/\log n) \leq K \log\log n$ to (2.1), the dangerous sets contribute

$$O\big(n/(\log\log n)^2\big) \cdot K \log\log n = o(n)$$

in total. As $\lambda_i = cn^{1/2}$ for the other sets and $m = n$, their total contribution is also at most $n/10$ for c large enough. ∎

3.2. The Lovett–Meka algorithm

Lovett and Mekka [45] substantially simplified the random-walk approach and extended it to give the following algorithmic version of the general partial coloring lemma.

Theorem 3.5. *Given an input matrix $A \in \mathbb{R}^{m\times n}$ and some fractional coloring $x_0 \in [-1,1]^n$ with k alive elements, for $i \in [m]$ let λ_i be such that*

$$\sum_i \exp\left(-\lambda_i^2/16\right) \le k/16. \tag{3.3}$$

Then there is a randomized polynomial time algorithm to find a coloring x with at most $k/2$ alive variables such that $|x(a_i) - x_0(a_i)| \le \lambda_i \|a_i\|_2$ for each row $i \in [m]$.

We remark that the colors produced by Theorem 3.5 lie in $[-1, 1]$, in contrast to $\{-1, 0, 1\}$ in Theorem 2.1, but this does not make any difference. Theorem 3.5 is also slightly stronger than Theorem 2.1 for $\lambda_i \ll 1$.

The key idea of the algorithm is that whenever a discrepancy constraint becomes tight or some variable reaches ± 1, one can simply do a random walk orthogonal to it. We now describe it formally. Without loss of generality, we assume that all variables are initially alive, i.e., $k = n$.

3.2.1. The algorithm

Let A_{t-1} be the set of alive variables at the beginning of time t. The algorithm maintains a linear subspace $V_{t-1} \subset \mathbb{R}^n$. Initially at $t = 1$, the coloring is x_0, $A_0 = [n]$ and $V_0 = \mathbb{R}^n$. The following is repeated for $\ell = O(\gamma^{-2})$ steps.

At time t, the algorithm chooses a random gaussian vector g_t in the subspace V_{t-1} and updates $x_t = x_{t-1} + \gamma g_t$, where γ is a small step size as usual.

(1) If $|x_t(j)| \ge 1$, set $V_t = V_{t-1} \cap e_j^{\perp}$, so that $x(j)$ will not be updated anymore.

(2) If $|x_t(a_i)| \ge \lambda_i \|a_i\|_2$, set $V_t = V_{t-1} \cap a_i^{\perp}$, so that row i incurs no further discrepancy.

Analysis. We assume that γ is small enough so that we can ignore the rounding error in the sketch below. By design, the algorithm ensures that $x_t(j) \in [-1, 1]$ for all j and that $x_t(a_i) \le \lambda_i \|a_i\|_2$ for all i. We now show that, with constant probability, at least half the variables reach ± 1.

For a linear subspace V, let $N(V)$ denote the standard multidimensional Gaussian distribution supported on V. By rotational invariance, a random vector $g \sim N(V)$ can be written as $g = g(1)v_1 + \cdots + g(d)v_d$ for some orthonormal basis $\{v_1, \ldots, v_d\}$ for V and $g(1), \ldots, g(d)$ iid $N(0, 1)$. We note the following fact.

Lemma 3.6. *Let V be a d-dimensional subspace of $\mathbb{R}^n$ and $g \sim N(V)$. Then for all $u \in \mathbb{R}^n$, $\langle g, u\rangle \sim N(0, \sigma^2)$ where $\sigma^2 \le \|u\|^2$. Moreover, for $i = 1, \ldots, n$ let σ_i be such that $\langle g, e_i\rangle \sim N(0, \sigma_i^2)$. Then $\sum_{i=1}^n \sigma_i^2 = d$.*

Proof. Let u' denote the projection of u onto V. Clearly, $\|u'\| \le \|u\|$. As $g \in V$, $\langle g, u\rangle = \langle g, u'\rangle$ and hence $\langle g, u\rangle \sim N(0, \|u'\|^2)$. For the second part, if $v_1, \ldots, v_d$ is an orthogonal basis for V, then $\sigma_i^2 = \sum_{j=1}^d \langle e_i, v_j\rangle^2$. Thus $\sum_{i=1}^n \sigma_i^2 = \sum_{i=1}^n \sum_{j=1}^d \langle e_i, v_j\rangle^2 = \sum_{j=1}^d \sum_{i=1}^n \langle v_j, e_i\rangle^2 = \sum_{j=1}^d \|v_j\|^2 = d$. ∎

Proof of Theorem 3.5. (Sketch) First we claim that in expectation, not many discrepancy constraints become tight in step (2) of the algorithm. This follows as for any time t, by Lemma 3.6 the discrepancy increment for each row a_i is distributed as $N(0, \leq \gamma^2 \|a_i\|^2)$. As $\ell = O(\gamma^{-2})$, by standard tail bounds $\Pr[|x_\ell(a_j) - x_0(a_j)| \geq \lambda_i \|a_i\|_2] = \exp(-\Omega(\lambda_i^2))$. As the λ_i satisfy (3.3), choosing the constants appropriately, the probability that more than $n/8$ discrepancy constraints becomes tight is at most $1/8$.

Let us condition on the above event. The proof now follows from a win–win argument. If more than $n/2$ elements reach ± 1, we are already done. If this does not happen, then at any time during the algorithm the subspace V_t has dimension at least $n - n/2 - n/8 \geq 3n/8$. By Lemma 3.6, as $\sum_j \mathbb{E}[\Delta x_t(j)^2] \geq (3n/8)\gamma^2$ and $\ell = O(\gamma^{-2})$ steps, the *energy* $\sum_j (x_\ell(j)^2 - x_0(j)^2)$ must increase by $\Omega(n)$ in expectation. But as $x_\ell(j)^2 - x_0(j)^2 \in [-1, 1]$ for all j, a simple argument can be used to show that at least $\Omega(n)$ variables reach ± 1 in expectation.

3.3. Direct approaches

The Lovett–Meka algorithm crucially uses the face structure of the polytope and does not seem to generalize to general convex bodies in the sense of Theorem 2.2. In particular, even if $\gamma_n(K) \geq 2^{-\delta n}$, condition (2.1) may not hold as it might require exponentially many facets to obtain any reasonable approximation of a general convex body K.

We now describe an extremely elegant and simple to state the algorithm due to Rothvoss [59], that finds a partial coloring in general convex bodies. We then describe a related algorithm by Eldan and Singh [28].

3.3.1. Rothvoss' algorithm

Let K be a symmetric convex body with $\gamma(K) \geq 2^{-\delta n}$. Take a random Gaussian $g \in \mathbb{R}^n$, and output the point closet to g in the body $K \cap [-1, 1]^n$, i.e., output

$$x^* = \operatorname{argmin}\{\|g - x\|_2 : x \in K \cap [-1, 1]^n\}.$$

That's it! The point x^* can be computed by a convex program, using a membership oracle for K.

Theorem 3.7 ([59]). *Let $\varepsilon > 0$ be a sufficiently small constant and $\delta := (3/2)\varepsilon \log_2(1/\varepsilon)$, and let K be a symmetric convex body with $\gamma_n(K) \geq \exp(-\delta n)$. Then whp, x^* has at least εn many coordinates ± 1.*

Analysis. The proof is also very elegant and uses Gaussian concentration for Lipschitz functions and the Sidak–Khatri lemma in a clever way.

The starting observation is that the distance $d(g, x^*)$ is at least $n^{1/2}/5$ with probability $1 - \exp(-\Omega(n))$. This follows as $x^* \in [-1, 1]^n$ and as $g(j) \sim N(0, 1)$ for each coordinate j, we have $\Pr[|g(j)| \geq 2] \geq 1/25$. On the other hand, $d(g, K) \leq 3(\delta n)^{1/2}$ with probability $1 - \exp(-\Omega(n))$ by Gaussian concentration for Lipschitz functions as $\gamma_n(K) \geq \exp(-\delta n)$.

Now, suppose for the sake of contradiction that fewer than εn coordinates of x^* are ± 1 for some g. Let I be the set of these coordinates. A key observation is that if x^* is an optimum solution to some convex program, and some constraint is not tight at x^*, then x^* remains optimum even when this constraint is removed. So x^* would still be the optimum solution, if replace $K \cap [-1,1]^n$ in the convex program by $K \cap S(I)$, where $S(I) = \{x : |x(j)| \le 1, j \in I\}$ is the intersection of the slabs corresponding to coordinates in I.

By the Sidak–Khatri lemma, $\gamma_n(K \cap S(I)) \ge \gamma_n(K)\gamma_n(S(I)) \ge \exp(-(\epsilon + 2\delta)n)$, and hence by Gaussian concentration the distance $d(g, x^*) = d(g, K \cap S(I)) \le 6(\varepsilon + \delta)n^{1/2}$ with probability $1 - \exp(-\Omega(n))$. So even if after a union bound over the $\approx \exp(\delta \ln(1/\delta)n)$ possible choices for I, one has $d(x^*, g) = O((\varepsilon + \delta)n^{1/2})$ whp. This contradicts the first observation that $d(x^*, g) \ge n^{1/2}/5$ whp.

3.3.2. Eldan–Singh algorithm

This algorithm is as simple to state and only requires linear optimization: Pick a random direction $c \in R^n$ and optimize over $K \cap [-1,1]^n$, i.e., output

$$x^* = \operatorname{argmax}\{c^T x : x \in K \cap [-1,1]^n\}.$$

Eldan and Singh [28] showed a result similar to Theorem 3.7. That is, for any $\epsilon > 0$ small enough, there is a $\delta > 0$ such that if $\gamma_n(K) \ge 2^{-\delta n}$ then whp x^* has at least εn coordinates ± 1 with constant probability.

4. ALGORITHMIC VERSION OF BANASZCZYKS'S RESULT

We now consider the algorithmic approaches for Banaszczyk's method. The first progress was by Bansal, Dadush, and Garg [8], who gave an efficient SDP-based algorithm to find an $O((\log n)^{1/2})$ discrepancy coloring for the Komlós problem. A deterministic algorithm for the problem was subsequently obtained by Levy, Ramadas, and Rothvoss [42].

Later, Bansal, Dadush, Garg, and Lovett [9] gave an algorithm for the general case of Banaszczyk's theorem with arbitrary convex body K. Their algorithm, called the Gram–Schmidt walk, combines linear algebra and random walks. Recently, Harshaw et al. [35] gave an optimal analysis of this walk.

We describe both these approaches below. We mention that finding an efficient algorithm for the prefix version of Banaszczyk's problem in Theorem 2.4 is still open.

Problem 4.1. Find an efficient algorithm for the prefix version of Banaszczyk's theorem. The case of prefix Komlós (Problem 1.4) would already be very interesting.

4.1. The Komlós problem

We describe the following result of Bansal, Dadush, and Garg [8].

Theorem 4.2 ([8]). *Given vectors $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_j\|_2 \le 1$ for $j \in [n]$, there is a polynomial time algorithm that finds an $O((\log n)^{1/2})$ discrepancy coloring whp.*

The algorithm is based on SDPs and is similar to that in Section 3.1, but it adds some extra constraints to the SDP so that the resulting solution has some additional desirable properties. To understand these properties, it is instructive to see what can go wrong with the partial coloring approach. To focus on the main ideas we consider the special case of Beck–Fiala problem, with the goal of finding an $O((d \log n)^{1/2})$ discrepancy coloring.

Recall that the $O(d^{1/2} \log n)$ bound using partial coloring was obtained by requiring zero discrepancy for large sets S, say of size $> 10d$. For small sets of size $s \leq 10d$, we can set the bound roughly $O(s^{1/2})$ (in Section 2.2.1 we used the bound $d^{1/2}$, but $O(s^{1/2} \ln(20d/s))$ also works). So as long as a set is large, it incurs zero discrepancy, and once it is small it incurs at most $O(d^{1/2})$ discrepancy in each partial coloring step.

The ideal process. Ideally, one would expect that once a set S becomes small, then whenever a constant fraction of the elements get colored colored globally in a partial coloring step, the size of S should also decrease geometrically. If so, this would actually give an $O(d^{1/2})$ discrepancy. However, the problem is that partial coloring does not give much control on which elements get colored, e.g., sets can incur discrepancy $O(d^{1/2})$ even if only $O(d^{1/2})$ of their elements get colored. This imbalance between the discrepancy and the *progress* a set makes in getting colored is the main barrier to improving the $O(d^{1/2} \log n)$ bound.

A concrete bad example. To see this more explicitly, let us consider the Lovett–Meka algorithm. Suppose the subspace V_{t-1} at time t is spanned by the orthonormal basis b, $e_{d+1}, \ldots, e_n$ where $b = d^{-1/2}(e_1 + \cdots + e_d)$. Then any update $\Delta x_t \in V_{t-1}$ has $\Delta x_t(1) = \cdots = \Delta x_t(d)$, and for the set $S = \{1, \ldots, d\}$, all variables get updated by the same amount, so if it incurs discrepancy $d^{1/2}$, the coloring progress is only $d^{1/2}$. In contrast, if the $\Delta x_t(1), \ldots, \Delta x_t(d)$ were independent, $\Omega(d)$ elements would get colored in expectation while incurring a discrepancy of $d^{1/2}$.

The key idea behind the algorithm of [8] is to ensure that even though the update Δx_t lies in some subspace that we cannot control, the coordinates $\Delta x_t(j)$ behave roughly independently in the sense that

$$\mathbb{E}\left[\left(\sum_j b(j)\Delta x_t(j)\right)^2\right] \leq \eta\left(\sum_j b(j)^2 \mathbb{E}\left[\Delta x_t(j)^2\right]\right) \quad \forall b \in \mathbb{R}^n, \tag{4.1}$$

where $\eta \geq 1$ is some fixed constant. Notice that if the $\Delta x_t(j)$ were independent or even pairwise independent, then (4.1) would be an equality with $\eta = 1$.

The algorithm will add an additional SDP constraint to ensure property (4.1). We describe this below and then give a sketch of the analysis.

4.1.1. Algorithm

Let (U, C) be the input set system. As usual, the algorithm starts with the coloring $x_0 = 0^n$. Let x_{t-1}, A_{t-1} denote the coloring and the set of alive variables at the beginning of t. Call a set $S \in C$ large if $|S \cap A_{t-1}| \geq 10d$.

Repeat the following for $t = 1, 2, \ldots, \ell$ until $A_\ell = \emptyset$.

(1) Solve the following SDP:

$$\left\| \sum_{j \in S} w_j \right\|^2 = 0 \quad \text{for all large } S, \tag{4.2}$$

$$\left\| \sum_j b(j) w_j \right\|^2 \leq 2 \sum_j b(j)^2 \|w_j\|^2 \quad \forall b \in \mathbb{R}^n, \tag{4.3}$$

$$\|w_j\|^2 \leq 1 \quad \text{for } j \in A_{t-1}, \text{ and else } \|w_j\|^2 = 0, \tag{4.4}$$

$$\sum_j \|w_j\|^2 \geq |A_{t-1}|/4. \tag{4.5}$$

(2) Let $\Delta x_t(j) = \gamma \langle g, v_i \rangle$ where g is a random Gaussian vector. Set $x_t = x_{t-1} + \Delta x_t$, and update A_t accordingly.

The infinitely many constraints (4.3) can be written compactly as $X \preceq 2\,\mathrm{diag}(X)$, where X is the Gram matrix with $X_{ij} = \langle w_i, w_j \rangle$.

4.1.2. Analysis

The constraints (4.2) ensures that $\Delta x_t(S) = 0$ for large sets, which are at most $|A_{t-1}|/10$ in number. The constraints (4.3) imply the property (4.1). The feasibility of the SDP follows from the following geometric result.

Theorem 4.3 ([10]). *Let $G \subset \mathbb{R}^n$ be an arbitrary subspace with dimension $\dim(G) = \delta n$. For any $\zeta > 0$ and $\eta > 1$ with $1/\eta + \zeta \leq \delta$, there is a $n \times n$ PSD matrix X satisfying:*

(i) $\langle hh^T, X \rangle = 0$ *for all $h \in G^\perp$, where $G^\perp$ is the subspace orthogonal to G.*

(ii) $X_{ii} \leq 1$ *for all $i \in [n]$.*

(iii) *The trace* $\mathrm{tr}(X) \geq \zeta n$.

(iv) $X \preceq \eta\,\mathrm{diag}(X)$.

In particular, choosing G to be the subspace orthogonal to all large rows and setting $\delta = 0.9$, $\eta = 2$, and $\zeta = 0.1$, Theorem 4.3 implies that the SDP is always feasible.

This algorithm can be viewed as an interesting extension of iterated-rounding, where the update lies in a subspace, and yet has interesting random-like properties.

Let us see why this helps. At any time t, the discrepancy for set S has Gaussian increments with variance $\mathbb{E}[(\sum_{j \in S} \Delta x_t(j))^2]$, which by (4.1) is at most $2 \sum_{j \in S} \mathbb{E}[\Delta x_t(j)^2]$, i.e., twice the variance injected into the elements of S. We will show that

$$\sum_t \left(\sum_{j \in S} \Delta x_t(j)^2 \right) = O(d)$$

whp, and hence the discrepancy of S will be a Gaussian with standard deviation $O(d^{1/2})$. A union bound over the sets then gives the desired $O((d \log n)^{1/2})$ bound.

To this end, let us define $\sum_{j \in S} x_t(j)^2$ as the energy of S at time t. By (4.2), any S incurs discrepancy only after it becomes small, and so from that time onward its energy

can increase by at most $O(d)$. A priori there is no reason why the total increase in energy of S should be related to $\sum_t \sum_{j\in S} \mathbb{E}[\Delta x_t(j)^2]$ (the total variance injected into the elements of S). For example, even for a single variable j if $x_t(j)$ fluctuates a lot over time, $\sum_t \Delta x_t(j)^2$ could be arbitrarily large, while the final energy is ≤ 1. More precisely, the change in energy of S at time t is

$$\sum_{j\in S}\left(x_t(j)^2 - x_{t-1}(j)^2\right) = \underbrace{2\sum_{j\in S} x_{t-1}(j)\Delta x_t(j)}_{I} + \underbrace{\sum_{j\in S}\Delta x_t(j)^2}_{II}.$$

Summing up over t, the left-hand side telescopes and equals the total increase in energy of S. But $\sum_t \Delta x_t(j)^2$ can be much larger than this if the sum of term I over time is very negative. However, constraint (4.1) turns out to be very useful again. In particular, term I is a mean-zero update, and by (4.1) its variance can be bounded as

$$\mathbb{E}\left[\left(\sum_{j\in S} x_{t-1}(j)\Delta x_t(j)\right)^2\right] \leq 2\sum_{j\in S} x_{t-1}(j)^2\Delta x_t(j)^2 \leq 2\sum_{j\in S}\Delta x_t(j)^2.$$

This implies that the contribution of I is quite small compared to $\sum_t \sum_{j\in S} \mathbb{E}[\Delta x_t(j)^2]$. A clean exposition based on supermartingale concentration is in [7].

4.2. The general setting

We now describe the algorithmic version of Theorem 2.3. For simplicity, we will assume that K is symmetric. This is almost without loss of generality, because if K is asymmetric with $\gamma_m(K) \geq 3/4$, then $K \cap -K$ is symmetric and $\gamma_m(K \cap -K) \geq 1/2$.

An immediate issue with making Theorem 2.3 algorithmic is that any explicit description of K to a reasonable accuracy could already require exponential space. A crucial first step was by Dadush, Garg, Nikolov, and Lovett [24] who reformulated Theorem 2.3 without any reference to K. To state this result, recall that a random vector $Y \in \mathbb{R}^m$ is σ-sub-Gaussian if for all test directions $\theta \in \mathbb{R}^m$,

$$\mathbb{E}\left[e^{\langle\theta, Y\rangle}\right] \leq e^{\sigma^2\|\theta\|_2^2/2}.$$

Roughly, this means that $\langle Y, \theta\rangle$ looks like a Gaussian random variable with variance at most σ^2 for every unit vector θ. Simplifying slightly to symmetric K, [24] showed the following.

Theorem 4.4 ([24]). *For any symmetric convex body K, Theorem 2.3 (up to the exact value of c) is equivalent to the following: Let $v_1, \ldots, v_n \in \mathbb{R}^m$ be vectors with $|v_j\|_2 \leq 1$. Then there exists a distribution D on colorings $\{-1, 1\}^n$, such that for x sampled from D, the random vector $\sum_{j=1}^n x(j)v_j$ is σ-sub-Gaussian for some $\sigma = O(1)$.*

Moreover, to get a constructive version of Theorem 2.3 for any K, it suffices to give an algorithm that can efficiently sample a coloring from D.

The idea behind Theorem 4.4 is that as $\gamma_m(K) \geq 1/2$, a random Gaussian $g \in \mathbb{R}^m$ satisfies $\Pr[g \in K] \geq 1/2$, or equivalently, $\Pr[\|g\|_K \leq 1] \geq 1/2$ where $\|\cdot\|_K$ is the norm with K as its unit ball. By standard tail bounds, this gives $\mathbb{E}[\|g\|_K] = O(1)$. The following result of Talagrand [66], together with Markov's inequality, directly gives Theorem 4.4.

Theorem 4.5 ([66]). *Let $K \subset \mathbb{R}^m$ be a symmetric convex body and $Y \in \mathbb{R}^m$ be a σ-sub-Gaussian random vector. Then for the standard Gaussian $g \in \mathbb{R}^m$,*

$$\mathbb{E}\big[\|Y\|_K\big] \le O(\sigma) \cdot \mathbb{E}\big[\|g\|_K\big].$$

Bansal, Dadush, Garg, and Lovett [9] designed an algorithm called the Gram–Schmidt walk (GS-walk), with the following guarantee.

Theorem 4.6 ([9]). *Given vectors $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_j\|_2 \le 1$, GS-walk outputs a coloring $x \in \{-1, 1\}^n$ such that $\sum_{j=1}^n x(j)v_j$ is sub-Gaussian with $\sigma \approx 6.32$.*

Harshaw, Sävje, Spielman, and Zhang [35] gave an improved analysis of the algorithm and showed that $\sigma = 1$, which is the best possible.

4.2.1. Gram–Schmidt walk algorithm

Before we describe the algorithm, we give some intuition. Suppose first that the vectors $v_1, \dots, v_n$ are orthogonal. Then, in fact a random coloring suffices. This follows as for any $\theta \in \mathbb{R}^m$, we have $\langle \theta, \sum_j x(j)v_j \rangle = \sum_j x(j)\langle \theta, v_j \rangle$, which for a random ± 1 coloring x is distributed as a sub-Gaussian with variance $\sum_j \langle \theta, v_j \rangle^2$, which is at most $\|\theta\|_2^2$ as the v_j are orthogonal and have at most unit length.

On the other extreme, suppose that $v_1, \dots, v_n$ are all identical and equal to some unit vector v. Then a random coloring is very bad and has variance n (instead of $O(1)$) in the direction $\theta = v$. The right thing here, of course, is to pair up the signs of $x(j)$. The general algorithm will handle these two extreme examples in a unified way, by trying to exploit the linear dependencies as much as possible while also using randomness.

We now describe the algorithm formally.

The Gram–Schmidt walk. Let $v_1, \dots, v_n$ be the input vectors. Let x_{t-1}, A_{t-1} denote the coloring and the set of alive elements at the beginning of time t.

Let $n(t) \in A_{t-1}$ be the largest indexed element alive at time t. This is called the *pivot* at time t and will play a special role. Let W_t be subspace spanned by the vectors in $A_{t-1} \setminus \{n(t)\}$ (i.e., all vectors alive at time t except $n(t)$). Let $v^\perp(t)$ be the orthogonal projection of the pivot $v_{n(t)}$ on $W_t^\perp$.

The algorithms works as follows. Initialize $x_0 = (0, \dots, 0)$ and $A_0 = [n]$.

At $t = 1, \dots, n$, do the following:

(1) Compute the update direction $u_t = (u_t(1), \dots, u_t(n)) \in \mathbb{R}^n$ as follows. Set $u_t(j) = 1$ for the pivot $j = n(t)$ and $u_t(j) = 0$ for $j \notin A_{t-1}$.
The $u_t(j)$ for the remaining $j \in A_{t-1} \setminus \{n(t)\}$ are defined by writing

$$v^\perp(t) = v_{n(t)} + \sum_{j \in A_{t-1} \setminus \{n(t)\}} u_t(j)v_j.$$

(2) Let $\delta_t^- < 0 < \delta_t^+$ be the unique negative and positive solutions for δ, respectively, to $\max_{j \in A_{t-1}} |x_{t-1}(j) + \delta u_t(j)| = 1$. Let

$$\delta_t = \begin{cases} \delta_t^- & \text{with probability } \delta_t^+/(\delta_t^+ - \delta_t^-), \\ \delta_t^+ & \text{with probability } -\delta_t^-/(\delta_t^+ - \delta_t^-). \end{cases}$$

(3) Update x_{t-1} randomly as $x_t = x_{t-1} + \delta_t u_t$. Update A_t accordingly.

Remark. Let us first see what the algorithm does for the two cases mentioned above. If the v_i are orthonormal, then $v^\perp(t) = v_{n(t)}$ as $v_{n(t}$ is orthogonal to W_t, and the algorithm only updates the color of the pivot. Moreover, at each time t $x(n(t))$ is set independently to ± 1, and so the algorithm eventually produces a completely random coloring. On the other hand, in the case where the v_i are identical, at each step t, as long as $n_t \geq 2$, the algorithm will exactly pair up the color of the pivot with the alive vector with the lowest index, resulting in overall discrepancy of at most 1.

Sketch of analysis. At each step, at least one element reaches -1 or 1, so the algorithm terminates in at most n steps.

Fix a vector $\theta \in \mathbb{R}^m$ with respect to which we want to show sub-Gaussianity of the discrepancy vector. Let $Y_t := \sum_{i=1}^n x_t(i) v_i$ and let $\text{disc}_t = \langle \theta, Y_t \rangle$. The goal is to show that

$$\mathbb{E}\big[e^{\text{disc}_n}\big] \leq e^{(\sigma^2/2)\|\theta\|_2^2}, \quad \text{for } \sigma = O(1).$$

Let us denote $\Delta x_t := x_t - x_{t-1} = \delta_t u_t$ and $\Delta \text{ disc}_t := \text{disc}_t - \text{disc}_{t-1}$. A key observation is that as u_t is chosen to satisfy $v^\perp(t) = \sum_{i=1}^n u_t(i) v_i$, we have

$$\Delta \text{ disc}_t = \sum_{i=1}^n \langle \theta, v_i \rangle \Delta x_t(i) = \delta_t \sum_{i=1}^n \langle \theta, v_i \rangle u_t(i) = \delta_t \big\langle \theta, v^\perp(t) \big\rangle \tag{4.6}$$

and hence depends only on the vector $v^\perp(t)$.

Proving sub-Gaussianity. We sketch the main idea. Let us first make a simplifying assumption that at each time t, the element to reach ± 1 is the pivot. So the elements get colored in the order $n, n-1, \ldots, 1$ and the pivot at time t is $n(t) = n - t + 1$. Let $w_1, \ldots, w_n$ be the orthonormal vectors obtained by applying the Gram–Schmidt orthonormalization procedure (GS) on the vectors $v_1, \ldots, v_n$ in that order. That is, $w_1 = v_1/\|v_1\|$ and for $i > 1$, w_i is the projection of v_i orthogonal to $v_1, \ldots, v_{i-1}$, normalized to have unit norm. Then $v^\perp(t) = \langle v_{n(t)}, w_{n(t)} \rangle w_{n(t)}$.

By (4.6), the overall discrepancy along θ is $\text{disc}_n(\theta) = \sum_{t=1}^n \delta_t \langle \theta, v^\perp(t) \rangle$. As δ_t is a mean-zero random variable chosen independently at time t, and $|\delta_t| \leq 2$, we have

$$\mathbb{E}\big[e^{\text{disc}_n(\theta)}\big] = \mathbb{E}\big[e^{\sum_{t=1}^n \delta_t \langle \theta, v^\perp(t) \rangle}\big] \leq e^{O(1) \cdot \sum_{t=1}^n \langle \theta, v^\perp(t) \rangle^2}.$$

But this is at most $e^{O(1) \cdot \|\theta\|_2^2}$, as desired, because

$$\sum_t \big\langle \theta, v^\perp(t) \big\rangle^2 = \sum_t \big\langle \theta, \langle v_{n(t)}, w_{n(t)} \rangle w_{n(t)} \big\rangle^2 \leq \sum_t \langle \theta, w_{n(t)} \rangle^2 \leq \|\theta\|_2^2,$$

as $|\langle v_{n(t)}, w_{n(t)}\rangle| \le \|w_{n(t)}\|_2 \|v_{n(t)}\|_2 \le 1$, and $\sum_i \langle \theta, w_i\rangle^2 \le \|\theta\|_2^2$ as the w_i are orthonormal.

In general the analysis needs some more care as non-pivot elements will also get colored during the process. But, roughly speaking, this only improves the bounds. If some non-pivot element x_k is colored at some time t, then the GS procedure (without v_k) will produce a different set of orthonormal vectors $\{w'_i\}$, but the increase in $\langle \theta, w'_{n(t)}\rangle^2 - \langle \theta, w_{n(t)}\rangle^2$, can be charged against the fact that k will never be a pivot anymore in the the future. We refer to [9] for the formal analysis.

5. APPROXIMATING HEREDITARY DISCREPANCY

In the previous sections we obtained bounds on the discrepancy of various classes of set systems and matrices. One can ask whether given a particular matrix A, can we efficiently determine $\mathrm{disc}(A)$. However, as described earlier in Theorem 3.2, discrepancy is hard to approximate in a very strong sense [21]. Intuitively, this is because discrepancy can be quite brittle, e.g., consider some matrix A with large discrepancy; however, if we duplicate each column of A, the resulting matrix has discrepancy 0.

Even though discrepancy is hard to approximate, in a surprising and remarkable result Matoušek, Nikolov, and Talwar [50] showed that $\mathrm{herdisc}(A)$ can be well approximated. Note that a priori it is not even clear how to certify (even approximately) that $\mathrm{herdisc}(A) \le k$, as it is the maximum over exponentially many quantities that themselves cannot be certified.

Theorem 5.1 ([50]). *There is an $O(\log m)^{3/2}$) approximation algorithm for computing the hereditary discrepancy of any $A \in \mathbb{R}^{m\times n}$.*

This is based on relating the hereditary discrepancy of a matrix to its γ_2-norm.

The γ_2-norm. For a matrix A, let $r(A) = \max_i (\sum_j A_{ij}^2)^{1/2}$ and $C(A) = \max_j (\sum_i A_{ij}^2)^{1/2}$ denote the largest ℓ_2-norm of rows and columns A. The $\gamma_2(A)$-norm of A is defined as

$$\gamma_2(A) = \min\{r(U)c(V) : UV = A\},$$

the smallest product $r(U)c(V)$ over all possible factorizations of A.

The quantity $\gamma_2(A)$ is efficiently computable using an SDP as follows. Consider vectors $w_1, \ldots, w_m$ corresponding to rows of U and $w_{m+1}, \ldots, w_{m+n}$ to columns of V. As $\alpha U, V/\alpha$ is also a valid factorization for any $\alpha > 0$, we can assume that $r(U) = c(V)$. Then, it is easily seen that $\gamma_2(A)$ is the smallest value t for which the following SDP is feasible.

$$\langle w_i, w_{j+m}\rangle = A_{ij} \quad \forall i \in [m], j \in [n] \text{ and } \langle w_i, w_i\rangle \le t \ \forall i \in [m+n]. \tag{5.1}$$

Theorem 5.1 follows from the following two facts.

Lemma 5.2. *For any $A \in \mathbb{R}^{m\times n}$ and factorization $A = UV$ with U, V arbitrary, we have that* $\mathrm{disc}(A) \le O(r(U)c(V)(\log 2m)^{1/2})$. *In particular,* $\mathrm{disc}(A) \le O(\gamma_2(A)(\log 2m)^{1/2})$.

This also implies that $\mathrm{herdisc}(A) \le O(\gamma_2(A)(\log 2m)^{1/2})$ as $\gamma_2(\cdot)$ itself is a hereditary function. Indeed, for any subset of columns S, we have $\gamma_2(A_{|S}) \le \gamma_2(A)$ as $A_{|S} = UV_{|S}$

and $C(V_{|S}) \leq C(V)$. The proof of Lemma 5.2 uses Banaszczyk's theorem in an interesting way.

Proof. Define the body $K = \{y : \|Uy\|_\infty \leq 2r(U)(\log 2m)^{1/2}\}$. Then $\gamma(K) \geq 1/2$ because for a random gaussian $g \sim N(0, I)$, $\Pr_g[\|Ug\|_\infty \geq 2r(U)(\log 2m)^{1/2}] \leq 1/2$.

As the columns of V have length at most $c(V)$ and $\gamma(K) \geq 1/2$, by Theorem 2.3 there exists $x \in \{-1, 1\}^n$ such that $y := Vx \in 5c(V)K$. By definition of K, this gives $\|Uy\|_\infty \leq 10r(U)c(V)(\log 2m)^{1/2}$, and as $Ax = Uy$, the result follows. ∎

Lemma 5.3. *For any $A \in \mathbb{R}^{m\times n}$, we have* $\text{herdisc}(A) \geq \Omega(\gamma_2(A)/\log m)$.

The proof of Lemma 5.3 establishes an interesting connection between the γ_2-norm and the determinant lower bound defined as follows.

$$\text{detlb}(A) = \max_k \max_{S\subset[m],T\subset[n],|S|=|T|=k} \left|\det(A_{S,T})\right|^{1/k},$$

where $A_{S,T}$ is the submatrix of A restricted to row and columns in S and T.

In a classical result, Lovász, Spencer, and Vesztergombi [44] showed that $\text{herdisc}(A) \geq \text{detlb}(A)/2$ for any matrix A. using a geometric view of hereditary discrepancy similar to that in Observation 2.1.1. In the other direction, Matoušek [48] showed that $\text{herdisc}(A) \leq O(\log(mn)(\log n)^{1/2}\,\text{detlb}(A))$. Interestingly, Matoušek's proof used Theorem 3.1 and duality for the SDP considered in Section 3.1. In particular, if the vector discrepancy is large for some subset of columns, there there must exist a sub-matrix with large detlb. This result was improved recently by Jiang and Reis [39] to $\text{herdisc}(A) \leq O((\log m \log n)^{1/2}\,\text{detlb}(A))$, and this bound is the best possible.

To prove Lemma 5.3, [50] show that $\text{detlb}(A) \geq \gamma_2(A)/\log m$ using the duality of the SDP (5.1) together with ideas of Matoušek [48].

The bounds in both Lemmas 5.2 and 5.3 are the best possible. However, the following conjecture seems quite plausible.

Conjecture 5.3.1. *There is an $O(\log mn)$ approximation algorithm for computing the hereditary discrepancy of any matrix A.*

As $\text{detlb}(A)$ and $\text{herdisc}(A)$ are within an $O(\log mn)$ factor, by the results of [44] and [39], one possible way to prove Conjecture 5.3.1 would be to give an $O(1)$ approximation for computing $\text{detlb}(A)$.

6. OTHER RECENT DIRECTIONS

We now discuss some other recent directions. First, we consider an interesting line of work on understanding the discrepancy of random instances. Next, we consider some results in the online setting where the vectors v_j are revealed over time and the sign $x(j)$ must be chosen immediately and irrevocably when v_j is revealed. Finally, we consider some matrix discrepancy problems, where one considers signed sums of matrices, instead of signed sums of vectors.

6.1. Random instances

In this survey, we restrict our attention to the work on the Beck–Fiala problem [2, 14, 29, 37, 57]. There are two natural probabilistic models here. Either each column has a 1 in exactly k positions chosen randomly out of the m choices, or the Bernoulli ensemble where each entry is 1 with probability $p = k/m$. The latter is slightly easier due to the lack of dependencies. For both these settings, an $O(k^{1/2})$ discrepancy can be achieved for the entire range of n and m, under fairly general conditions [14, 57]. These results are also algorithmic.

An interesting recent line of work shows that in fact much smaller discrepancy is possible if $n \gg m$. Franks and Saks [31] showed that $\operatorname{disc}(A) \leq 2$ with high probability for a fairly general class of random matrices A if $n = \Omega(m^3 \log^2 m)$. Independently, Hoberg and Rothvoss [37] showed that $\operatorname{disc}(A) \leq 1$ whp for the Bernoulli ensemble if $n = \Omega(m^2 \log m)$, provided that $mp = \Omega(\log n)$. Both these results use Fourier based techniques and are non-algorithmic.

Let us note that $n = \Omega(m \log m)$ is necessary to achieve $O(1)$ discrepancy, provided that p is not too small. Indeed, if we fix any coloring x, and consider a random instance, the probability that a fixed row has discrepancy $O(1)$ is $O((pn)^{-1/2})$, so the probability that each row has discrepancy $O(1)$ is at most $(pn)^{-\Omega(m)}$. As there are (only) 2^n possible colorings, a first moment argument already requires that $2^n (pn)^{-m} = \Omega(1)$.

So a natural question is whether the discrepancy is actually $O(1)$ for $n = \Omega(m \log m)$. Curiously, the Fourier-based methods seem to require $n = \Omega(m^2)$ even for $p = 1/2$. However, subsequent results show this optimal dependence using the second moment method. Potukuchi [56] showed that $\operatorname{disc}(A) \leq 1$ if $n = \Omega(m \log m)$ for the dense case of $p = 1/2$. The sparse setting with $p \ll 1$ turns out to be more subtle, and was only recently resolved by Altschuler and Weed [2] using a more sophisticated approach based on the conditional second moment method together with Stein's method of exchangeable pairs. They show the following result.

Theorem 6.1 ([2]). *Let $A \in \{0,1\}^{m \times n}$ be a random matrix with each entry independently chosen to be 1 with probability $p := p(n)$. Then there is a constant $c > 0$ such that if $n \geq cm \log m$, then $\operatorname{disc}(A) \leq 1$ whp.*

The results of [2, 56] are also non-algorithmic, and given the use of the probabilistic method it seems unlikely that they can be made algorithmic. However, one may wonder if this can be done under weaker assumptions such as when $n \gg m^{10}$.

Problem 6.2. Is there an efficient algorithm to find a coloring with expected discrepancy $O(1)$ for random instances of the Beck–Fiala problem when $n = m^{\Omega(1)}$.

Smoothed analysis. A substantial generalization of the random setting is the smoothed analysis setting, where the instance is obtained by taking underlying worst-case instance and perturbing it by a small random noise [65]. Recently, [12] studied the prefix-Komlós problem in this setting, where the vectors $v_1, \ldots, v_n$ are chosen adversarially and then v_j is perturbed by an independent random noise vector u_j.

Theorem 6.3 ([12]). *If the covariance* $\mathrm{Cov}(u_j) \succcurlyeq \varepsilon^2 I_m$ *for some* $\epsilon \geq 1/\mathrm{poly}(m, \log n)$. *Then, whp each prefix has discrepancy* $O((\log m + \log\log n)^{1/2})$.

This improves the dependence on n in Theorem 2.4 to doubly logarithmic, even if the noise is quite small, e.g., a vector is changed only with probability $1/\mathrm{poly}(m, \log n)$ in a single random coordinate. The techniques for random instances do not directly work here as these methods crucially use various special properties of random instances.

An interesting question is whether Theorem 6.1 can be extended to the smoothed setting.

Problem 6.4. Does the Beck–Fiala problem have $O(1)$ expected discrepancy in the smoothed setting, for a reasonably small noise rate, when $n = m^{\Omega(1)}$.

6.2. Online setting

In all the results considered thus far, we assumed that the vectors $v_1, \ldots, v_n \in \mathbb{R}^m$ are all given in advance. Another natural model is the online setting, first studied by Spencer [62], where the vector v_t is revealed at time t and a sign $x(t)$ must be chosen irrevocably without the knowledge of the vectors that will arrive in the future. The goal is to keep the discrepancy $\|d_t\|_\infty$ any time t as small as possible, where $d_t = x(1)v_1 + \cdots + x(t)v_t$ is the discrepancy at end of time t.

We restrict our focus here to the online Komlós setting. Notice that setting $x(t)$ randomly to ± 1 also works in the online setting, but this gives $\Omega(n^{1/2})$ dependence on n. Unfortunately, this dependence is unavoidable in general—at each step t an adversary can choose the vector v_t to be orthogonal to the current discrepancy vector d_{t-1} causing $\|d_t\|_2$ (and hence $\|d_t\|_\infty$) to grow as $\Omega(t^{1/2})$ with time. More refined lower bounds are also known [16, 64].

Interestingly, it turns out that the dependence on n can be substantially improved if the vectors v_t are chosen in a less adversarial manner.

Stochastic model. Here the vectors are chosen randomly and independently from some distribution D, that is known to the algorithm [11, 13, 15, 34]. For the Komlós setting, [11] showed the following.

Theorem 6.5 ([11]). *Let* D *be any distribution on unit vectors in* $\mathbb{R}^m$. *There is an online algorithm that given vectors sampled iid from* D, *achieves discrepancy* $O(\log^4 mn)$ *whp.*

These results are based on a greedy deterministic algorithms that choose the sign $x(t)$ based on a suitable potential function.

Let us consider the simpler setting of ℓ_2 discrepancy and where D is the uniform distribution over the unit sphere S^{m-1}. We sketch the proof of an $O(m^{1/2})$ bound (which is the best possible for ℓ_2-discrepancy even offline, e.g., for m orthonormal vectors).

Consider the potential $\Phi_t = \|d_t\|_2^2$. Upon given v_t, the algorithm chooses $x(t)$ to minimize the increase in $\Delta\Phi_t = \Phi_t - \Phi_{t-1}$. This evaluates to

$$\|d_{t-1} + x(t)v_t\|^2 - \|d_{t-1}\|_2^2 = 2x(t)\langle d_t, v_t\rangle + |v_t|_2^2 = 2x(t)\langle d_t, v_t\rangle + 1,$$

and hence setting $x(t) = -\text{sign}(\langle d_t, v_t\rangle)$ gives $\Delta\Phi_t = -2|\langle d_t, v_t\rangle| + 1$.

As v_t is uniform in S^{m-1}, we have that in expectation $\mathbb{E}_D|\langle d_{t-1}, v_t\rangle| \approx m^{-1/2}\|d_{t-1}\|_2$ for any d_{t-1}. This gives that $\mathbb{E}[\Delta\Phi_t] \ll 0$, and hence Φ_t has a strong negative drift whenever $|d_{t-1}| \gg m^{1/2}$. Using standard arguments, this implies that the discrepancy is $O(m^{1/2})$ whp at any given time.

The case of general distributions D is harder as $\mathbb{E}_D|\langle d_{t-1}, v_t\rangle|$ need not be large for every d_{t-1}. For example, if most of the probability mass of D lies in some subspace M, and d_{t-1} is orthogonal to M. However, one can still make this approach work by considering more complicated potential functions, that in addition to penalizing d_{t-1} with large norm, also penalize d_{t-1} if it gets close to certain undesirable regions.

Oblivious adversary model. Recently, these results were considered in the much more general *oblivious adversary* model. Here, the adversary knows the online algorithm and can pick the vectors accordingly, but it must choose them in advance before the online algorithm begins its execution. Equivalently, it cannot see the internal random choices made by the algorithm.

Notice that the oblivious setting generalizes both the stochastic setting and the worst case offline setting. Moreover, unlike for the stochastic model, here the $\Omega(n^{1/2})$ lower bound holds for any deterministic online algorithm, as d_{t-1} is completely determined by $v_1, \ldots, v_{t-1}$ and the adversary can always pick v_t orthogonal to d_{t-1}. So any nontrivial algorithm in this model must use its internal randomness cleverly.

In a recent breakthrough, Alweiss, Liu, and Sawhney [3] showed the following remarkable result.

Theorem 6.6 ([3]). *For any $\delta > 0$, vectors $v_1, v_2, \ldots, v_n \in \mathbb{R}^m$ with $\|v_t\|_2 \le 1$ for all $t \in [n]$, the algorithm maintains $\|d_t\|_\infty = O(\log(mn/\delta))$ for all $t \in [n]$ with probability $1 - \delta$.*

Choosing $\delta = 1/n^2$ gives that each prefix has discrepancy $O(\log mn)$ whp, almost matching the offline $O((\log mn)^{1/2})$ bound for prefix discrepancy given by Theorem 2.4. Moreover, the algorithm is extremely elegant and simple to describe.

Self-balancing walk algorithm. Let $c = 30 \log mn/\delta$. At each time t,

(1) If $|d_{t-1}|_\infty > c$ or if $|\langle d_{t-1}, v_t\rangle| > c$, declare failure.

(2) Set $x_t = 1$ with probability $1/2 - \langle d_{t-1}, v_t\rangle/2c$ and $x_t = -1$ otherwise.

The algorithm can be viewed as a randomized version of the greedy algorithm that picks the sign randomly if v_t and d_{t-1} are orthogonal, and otherwise uses the correlation between them to create a bias to move d_t closer to the origin.

The proof is a based on a clever stochastic domination argument and induction, and shows that as long as the algorithm does not declare failure, the distribution of d_t is less spread out than $N(0, 2\pi c I)$.

Theorem 6.6 is remarkable in many ways. First, it gives a simple linear time algorithm to obtain $O(\log n)$ discrepancy for the Komlós problem. Second, it also matches the

best algorithmic bound that we currently know for the prefix-Komlós problem in the offline setting. So any improvement of Theorem 6.6 would be extremely interesting.

Problem 6.7. Design an online algorithm for the Komlós problem, or for the prefix Komlós problem, that achieves an $O((\log mn)^{1/2})$ discrepancy.

For recent partial progress in this direction, see [43].

6.3. Matrix discrepancy

So far we only considered problems involving a signed sum of vectors. It is also very interesting to consider signed sums of more general objects such as matrices. An important problem of this type, with application to various fields, is the Kadison–Singer problem [40]. Below is an equivalent formulation in terms of discrepancy due to Weaver [68].

Kadison–Singer problem [68]. Let $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ be rank-1 Hermitian matrices satisfying $\sum_j A_j = I$ and $\|A_j\|_{\mathrm{op}} \le \delta$ for all $j \in [n]$, where $\delta \le 1/2$. Is there a ± 1 coloring x such that $\|\sum_j x(j)A_j\|_{\mathrm{op}} \le 1-\eta$, for some fixed constant $\eta > 0$ independent of n and d?

More generally, one can ask how small can the *discrepancy* $\|\sum_j x(j)A_j\|_{\mathrm{op}}$ be over all possible ± 1 colorings x. For a random coloring, standard matrix concentration results [55] give a bound of $O((\delta \log d)^{1/2})$, which does not give anything useful for the Kadison–Singer problem for large d. In a major breakthrough, Marcus, Spielman, and Srivastava [46] showed a bound of $O(\sqrt{\delta})$, without any dependence on d, using the method of interlacing polynomials. This bound is also the best possible [68]. These techniques are very different and we do not discus them here.

Their result however is non-constructive and obtaining an algorithmic version in an outstanding open question.

Problem 6.8. Is there an algorithmic version for the Kadison–Singer problem, even for the weaker bound of $1-\eta$ instead of $O(\sqrt{\delta})$.

Matrix Spencer problem. Another very interesting question, proposed originally by Raghu Meka, is the following matrix version of the Spencer's problem: given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{n\times n}$ with $\|A_j\|_{\mathrm{op}} \le 1$, find a ± 1 coloring x to minimize $\|\sum_j x(j)A_j\|_{\mathrm{op}}$.

Notice that if the A_j are diagonal, this is equivalent to Spencer's problem for $m = n$. Again, standard matrix concentration bounds imply a $O((n \log n)^{1/2})$ bound for random coloring, and the question is whether better bounds are possible.

Conjecture 6.8.1. *The matrix Spencer problem has discrepancy* $O(n^{1/2})$.

Very recently, Hopkins, Raghavendra, and Shetty [38] proved Conjecture 6.8.1 when the A_j have rank $n^{1/2}$, or, more generally, when $\|A_j\|_F \le n^{1/2}$. This result is based on an interesting new connection between discrepancy and communication complexity, and they also use this to give an alternate new proof of Spencer's result in classical setting. Another related result is due to Dadush, Jiang, and Reis [25].

FUNDING

This work was partially supported by the NWO VICI grant 639.023.812.

REFERENCES

[1] N. Alon and J. H. Spencer, *The probabilistic method*. John Wiley & Sons, 2016.

[2] D. J. Altschuler and J. Niles-Weed, The discrepancy of random rectangular matrices. 2021, arXiv:2101.04036.

[3] R. Alweiss, Y. P. Liu, and M. Sawhney, Discrepancy minimization via a self-balancing walk. In *Proceedings of STOC*, pp. 14–20, ACM, 2021.

[4] W. Banaszczyk, Balancing vectors and Gaussian measures of n-dimensional convex bodies. *Random Structures Algorithms* **12** (1998), no. 4, 351–360.

[5] W. Banaszczyk, On series of signed vectors and their rearrangements. *Random Structures Algorithms* **40** (2012), no. 3, 301–316.

[6] N. Bansal, Constructive Algorithms for Discrepancy Minimization. In *Proceedings of FOCS 2010*, pp. 3–10, IEEE, 2010.

[7] N. Bansal, On a generalization of iterated and randomized rounding. In *Symposium on theory of computing, STOC*, pp. 1125–1135, ACM, 2019.

[8] N. Bansal, D. Dadush, and S. Garg, An algorithm for Komlós conjecture matching Banaszczyk's bound. In *Proceedings of FOCS 2016*, pp. 788–799, IEEE, 2016.

[9] N. Bansal, D. Dadush, S. Garg, and S. Lovett, The Gram–Schmidt walk: a cure for the Banaszczyk blues. *Theory Comput.* **15** (2019), 1–27.

[10] N. Bansal and S. Garg, Algorithmic discrepancy beyond partial coloring. In *Proceedings of STOC 2017*, pp. 914–926, ACM, 2017.

[11] N. Bansal, H. Jiang, R. Meka, S. Singla, and M. Sinha, Online discrepancy minimization for stochastic arrivals. In *Proceedings of SODA*, pp. 2842–2861, ACM-SIAM, 2021.

[12] N. Bansal, H. Jiang, R. Meka, S. Singla, and M. Sinha, Prefix discrepancy, smoothed analysis, and combinatorial vector balancing. 2021, arXiv:2111.07049.

[13] N. Bansal, H. Jiang, S. Singla, and M. Sinha, Online vector balancing and geometric discrepancy. In *Proceedings of STOC*, pp. 1139–1152, ACM, 2020.

[14] N. Bansal and R. Meka, On the discrepancy of random low degree set systems. In *Proceedings of SODA 2019*, pp. 2557–2564, ACM-SIAM, 2019.

[15] N. Bansal and J. H. Spencer, On-line balancing of random inputs. *Random Structures Algorithms* **57** (2020), no. 4, 879–891.

[16] I. Bárány, On a Class of Balancing Games. *J. Combin. Theory Ser. A* **26** (1979), no. 2, 115–126.

[17] I. Bárány, On the power of linear dependencies. In *Building bridges*, pp. 31–45, Springer, 2008.

[18] J. Beck, Balanced two-colorings of finite sets in the square I. *Combinatorica* **1** (1981), no. 4, 327–335.

[19] J. Beck and T. Fiala, "Integer-making" theorems. *Discrete Appl. Math.* **3** (1981), no. 1, 1–8.

[20] J. Beck and V. Sós, Discrepancy theory. In *Handbook of combinatorics 2*, pp. 1405–1446, North-Holland 1995.

[21] M. Charikar, A. Newman, and A. Nikolov, Tight hardness results for minimizing discrepancy. In *Symposium on discrete algorithms, SODA*, pp. 1607–1614, ACM-SIAM, 2011.

[22] B. Chazelle, *The discrepancy method: randomness and complexity*. Cambridge University Press, 2001.

[23] W. Chen, A. Srivastav, G. Travaglini, et al., *A panorama of discrepancy theory*. Lecture Notes in Math. 2107, Springer, 2014.

[24] D. Dadush, S. Garg, S. Lovett, and A. Nikolov, Towards a constructive version of Banaszczyk's vector balancing theorem. In *Proceedings of APPROX/RANDOM*, pp. 28:1–28:12, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.

[25] D. Dadush, H. Jiang, and V. Reis, A new framework for matrix discrepancy: partial coloring bounds via mirror descent. 2021, arXiv:2111.03171.

[26] M. Drmota and R. F. Tichy, *Sequences, discrepancies and applications*. Springer, 1997.

[27] F. Eisenbrand, D. Pálvölgyi, and T. Rothvoß, Bin packing via discrepancy of permutations. *ACM Trans. Algorithms* **9** (2013), 24:1–24:15.

[28] R. Eldan and M. Singh, Efficient algorithms for discrepancy minimization in convex sets. *Random Structures Algorithms* **53** (2018), no. 2, 289–307.

[29] E. Ezra and S. Lovett, On the Beck–Fiala conjecture for random set systems. *Random Structures Algorithms* **54** (2019), no. 4, 665–675.

[30] C. Franks, A simplified disproof of Beck's three permutations conjecture and an application to root-mean-squared discrepancy. 2018, arXiv:1811.01102.

[31] C. Franks and M. Saks, On the discrepancy of random matrices with many columns. *Random Structures Algorithms* **57** (2020), no. 1, 64–96.

[32] A. A. Giannopoulos, On some vector balancing problems. *Studia Math.* **122** (1997), no. 3, 225–234.

[33] E. D. Gluskin, Extremal properties of orthogonal parallelepipeds and their applications to the geometry of Banach spaces. *Math. USSR, Sb.* **64** (1989), no. 1, 85.

[34] N. Haghtalab, T. Roughgarden, and A. Shetty, Smoothed analysis with adaptive adversaries. 2021, arXiv:2102.08446.

[35] C. Harshaw, F. Sävje, D. A. Spielman, and P. Zhang, Balancing covariates in randomized experiments using the Gram–Schmidt walk. 2019, arXiv:1911.03071.

[36] N. J. A. Harvey, R. Schwartz, and M. Singh, Discrepancy without partial colorings. In *Approx/random*, pp. 258–273, LIPIcs 28, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2014.

[37] R. Hoberg and T. Rothvoss, A Fourier-analytic approach for the discrepancy of random set systems. In *Proceedings of SODA*, pp. 2547–2556, ACM-SIAM, 2019.

[38] S. B. Hopkins, P. Raghavendra, and A. Shetty, Matrix discrepancy from quantum communication. 2021, arXiv:2110.10099v1.

[39] H. Jiang and V. Reis, A tighter relation between hereditary discrepancy and determinant lower bound. 2021, arXiv:2108.07945.

[40] R. V. Kadison and I. M. Singer, Extensions of pure states. *Amer. J. Math.* **81** (1959), 383–400.

[41] L.-C. Lau, R. Ravi, and M. Singh, *Iterative methods in combinatorial optimization*. Cambridge University Press, 2011.

[42] A. Levy, H. Ramadas, and T. Rothvoss, Deterministic discrepancy minimization via the multiplicative weight update method. In *Proceedings of IPCO 2017*, pp. 380–391, Springer, 2017.

[43] Y. P. Liu, A. Sah, and M. Sawhney, A gaussian fixed point random walk. 2021, arXiv:2104.07009.

[44] L. Lovász, J. Spencer, and K. Vesztergombi, Discrepancy of set-systems and matrices. *European J. Combin.* **7** (1986), no. 2, 151–160.

[45] S. Lovett and R. Meka, Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.* **44** (2015), no. 5, 1573–1582.

[46] A. W. Marcus, D. A. Spielman, and N. Srivastava, Interlacing families II: mixed characteristic polynomials and the Kadison–Singer problem. *Ann. of Math.* **182** (2015), 327–350.

[47] J. Matoušek, *Geometric discrepancy: an illustrated guide*. Algorithms Combin. 18, Springer, 2009.

[48] J. Matoušek, The determinant bound for discrepancy is almost tight. *Proc. Amer. Math. Soc.* **141** (2013), no. 2, 451–460.

[49] J. Matoušek and A. Nikolov, Combinatorial discrepancy for boxes via the γ_2 norm. In *Symposium on computational geometry, SoCG*, pp. 1–15, LIPIcs 34, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2015.

[50] J. Matoušek, A. Nikolov, and K. Talwar, Factorization norms and hereditary discrepancy. *Int. Math. Res. Not.* **2020** (2018), no. 3, 751–780.

[51] J. Matoušek and J. Spencer, Discrepancy in arithmetic progressions. *J. Amer. Math. Soc.* **9** (1996), no. 1, 195–204.

[52] A. Newman, O. Neiman, and A. Nikolov, Beck's three permutations conjecture: a counterexample and some consequences. In *Proceedings of FOCS 2012*, pp. 253–262, IEEE, 2012.

[53] A. Nikolov, *New computational aspects of discrepancy theory*. PhD thesis, Rutgers University, 2014.

[54] A. Nikolov, Tighter bounds for the discrepancy of boxes and polytopes. *Mathematika* **63** (2017), no. 3, 1091–1113.

[55] R. Oliveira, Sums of random hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab.* **15** (2010).

[56] A. Potukuchi, Discrepancy in random hypergraph models. 2018, arXiv:1811.01491.

[57] A. Potukuchi, A spectral bound on hypergraph discrepancy. In *Proceedings of ICALP*, pp. 93:1–93:14, LIPIcs 168, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2020.

[58] K. F. Roth, Remark concerning integer sequences. *Acta Arith.* **9** (1964), 257–260.

[59] T. Rothvoß, Constructive discrepancy minimization for convex sets. In *Proceedings of FOCS 2014*, pp. 140–145, IEEE, 2014.

[60] T. Rothvoss, Better bin packing approximations via discrepancy theory. *SIAM J. Comput.* **45** (2016), no. 3, 930–946.

[61] W. M. Schmidt, Irregularities of distribution vii. *Acta Arith.* **21** (1972), 45–50.

[62] J. Spencer, Balancing games. *J. Combin. Theory Ser. B* **23** (1977), no. 1, 68–74.

[63] J. Spencer, Six standard deviations suffice. *Trans. Amer. Math. Soc.* **289** (1985), no. 2, 679–706.

[64] J. Spencer, Balancing vectors in the max norm. *Combinatorica* **6** (1986), no. 1, 55–65.

[65] D. A. Spielman and S. Teng, Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Commun. ACM* **52** (2009), no. 10, 76–84.

[66] M. Talagrand, *The generic chaining: upper and lower bounds of stochastic processes*. Springer, 2005.

[67] T. Tao, The Erdős discrepancy problem. *Discrete Anal.* (2016).

[68] N. Weaver, The Kadison–Singer problem in discrepancy theory. *Discrete Math.* **278** (2004), 227–239.

[69] H. Weyl, Über die Gleichverteilung von Zahlen mod. Eins. *Math. Ann.* **77** (1916), 313–352.

NIKHIL BANSAL

Computer Science and Engineering (CSE), University of Michigan, Ann Arbor, MI 48109-2121, USA, bansaln@umich.edu