

Gaussian discrepancy: a probabilistic relaxation of vector balancing

Sinho Chewi
Patrik Gerber
Philippe Rigollet
MIT

SCHEWI@MIT.EDU
 PRGERBER@MIT.EDU
 RIGOLLET@MIT.EDU

Paxton Turner
Harvard University

PAXTONTURNER@FAS.HARVARD.EDU

Abstract

We introduce a novel relaxation of combinatorial discrepancy called *Gaussian discrepancy*, whereby binary signings are replaced with correlated standard Gaussian random variables. This relaxation effectively reformulates an optimization problem over the Boolean hypercube into one over the space of correlation matrices. We show that Gaussian discrepancy is a tighter relaxation than the previously studied vector and spherical discrepancy problems, and we construct a fast online algorithm that achieves a version of the Banaszczyk bound for Gaussian discrepancy. This work also raises new questions such as the Komlós conjecture for Gaussian discrepancy, which may shed light on classical discrepancy problems.

Keywords: combinatorial optimization, Gaussian discrepancy, online discrepancy, spherical discrepancy, vector discrepancy

1. INTRODUCTION AND OVERVIEW OF THE PAPER

In this work we introduce a probabilistic relaxation of the classical combinatorial discrepancy problem that we call *Gaussian discrepancy*. In this section, we first briefly survey discrepancy theory and formally define our relaxation. Then we discuss our main results which consist of (i) sharp comparisons of Gaussian discrepancy to previously studied relaxations, and (ii) a fast algorithm for online Gaussian discrepancy. We conclude this overview with some open problems and an outline of the remainder of the paper.

1.1. Discrepancy theory and Gaussian discrepancy

1.1.1. BACKGROUND ON COMBINATORIAL DISCREPANCY

Discrepancy theory is a rich area of mathematics which has both inspired the development of novel tools for its study, and found numerous applications in a variety of fields such as combinatorics, computer science, geometry, optimization, and statistics; see the textbooks Matoušek (1999); Chazelle (2000). In one of its most fundamental forms, discrepancy asks the following question: given an $m \times n$ matrix $A \in \mathbb{R}^{m \times n}$, determine the value of the *discrepancy* of A , defined as

$$\text{disc}(A) := \min_{\sigma \in \{\pm 1\}^n} \|A\sigma\|_\infty = \min_{\sigma \in \{\pm 1\}^n} \max_{1 \leq i \leq m} |(A\sigma)_i|. \quad (1)$$

This question can be interpreted in terms of *vector balancing*: if v_1, \dots, v_n denote the columns of A , then we are looking for a *signing*, that is, a vector of signs $\sigma \in \{\pm 1\}^n$, that makes the signed sum $\sum_{j=1}^n \sigma_j v_j$ have small entries.

A seminal result in this area, due to [Spencer \(1985\)](#) and independently [Gluskin \(1989\)](#) (see also [Giannopoulos \(1997\)](#)), states that

$$\text{disc}(A) \leq 6\sqrt{n} \quad (2)$$

when $m = n$ and the entries of A are bounded in magnitude by 1. This remarkable result is the best possible, up to the constant factor 6, and should be compared with the discrepancy incurred by a signing chosen uniformly at random which is of order $\Theta(\sqrt{n \log n})$. A far-reaching extension of this result is the *Komlós conjecture* ([Spencer, 1985](#)).

Conjecture 1 (Komlós conjecture) *There exists a constant $K > 0$ such that for any matrix A whose columns have Euclidean norm at most 1, it holds that*

$$\text{disc}(A) \leq K.$$

This conjecture remains one of the most important open problems in the field, and the best known bound for the Komlós problem, due to [Banaszczyk \(1998\)](#), yields $\text{disc}(A) = O(\sqrt{\log(m \wedge n)})$.¹ The Komlós conjecture contains as a special case the long-standing Beck–Fiala conjecture which states that if A has t -sparse columns in $\{0, 1\}^m$, then $\text{disc}(A) = O(\sqrt{t})$ ([Beck and Fiala, 1981](#)).

The original proofs of these results are non-constructive in the sense that they do not readily yield efficient algorithms for computing signings which achieve these discrepancy upper bounds. In the last decade, considerable effort was devoted to matching these upper bounds *algorithmically*. Starting with the breakthrough work of [Bansal \(2010\)](#), there are now a number of algorithmic results matching Spencer’s bound ([Lovett and Meka, 2012](#); [Harvey et al., 2014](#); [Levy et al., 2017](#); [Rothvoss, 2017](#); [Eldan and Singh, 2018](#)). The task of making Banaszczyk’s bound algorithmic was more challenging, and it was settled in the last few years in a line of works ([Bansal et al., 2016](#); [Levy et al., 2017](#); [Bansal et al., 2018](#); [Dadush et al., 2019](#)).

Recently, *online* discrepancy minimization ([Bansal and Spencer, 2020](#); [Bansal et al., 2020](#); [Alweiss et al., 2021](#); [Bansal et al., 2021](#); [Liu et al., 2022](#)) has seen increasing interest and has led to a new perspective on Banaszczyk’s result. In the *oblivious* online setting, an adversary picks in advance vectors $v_1, \dots, v_T \in \mathbb{R}^m$, each with Euclidean norm at most 1. During each round $t = 1, \dots, T$, the algorithm receives the vector v_t , and it must output a sign $\sigma_t \in \{\pm 1\}$. The goal of the algorithm is to minimize the maximum discrepancy incurred at any time, *i.e.* the quantity $\max_{t \in [T]} \|\sum_{s=1}^t \sigma_s v_s\|_\infty$. [Alweiss et al. \(2021\)](#) conjecture the following online version of the Banaszczyk bound.

Conjecture 2 (Online Banaszczyk) *There exists a randomized algorithm for online balancing in the oblivious adversarial setting that with high probability achieves the bound*

$$\max_{t \in [T]} \left\| \sum_{s=1}^t \sigma_s v_s \right\|_\infty \lesssim \sqrt{\log(mT)},$$

for any sequence $v_1, \dots, v_T \in \mathbb{R}^m$ of vectors with Euclidean norm at most 1.

This would be nearly optimal, as a lower bound of $\tilde{\Omega}(\sqrt{\log T})$ is established in [Bansal et al. \(2020\)](#).

1. Here $m \wedge n = \min\{m, n\}$.

1.1.2. GAUSSIAN DISCREPANCY AND THE COUPLING PERSPECTIVE

Motivated by the aforementioned longstanding conjectures, recent algorithmic progress, and the goal of shedding new light on the classical discrepancy minimization problem, in this work we introduce a novel relaxation called *Gaussian discrepancy*. Our route to Gaussian discrepancy is through an alternative perspective on the discrepancy objective (1) based on couplings of random variables.

Recall that a Rademacher random variable is distributed uniformly on $\{\pm 1\}$. A *coupling* of Rademacher random variables is a random vector for which each marginal distribution is Rademacher. Since a signing $\sigma \in \{\pm 1\}^n$ and its negative $-\sigma$ achieve the same discrepancy, the uniform distribution on the set of optimal signings furnishes an *optimal* Rademacher coupling that minimizes the right-hand-side of (1) in expectation. Precisely, it holds that

$$\text{disc}(A) = \min \left\{ \mathbb{E} \|A\sigma\|_\infty : \mathbb{P}(\sigma_j = -1) = \mathbb{P}(\sigma_j = +1) = \frac{1}{2} \text{ for all } j \in [n] \right\}, \quad (3)$$

where the minimization above is over couplings of Rademacher random variables.²

The coupling perspective plays an important role in discrepancy theory and its applications. The recent algorithmic proof of Banaszczyk's theorem by [Bansal et al. \(2018\)](#) relies on the equivalence between Banaszczyk's theorem³ and the existence of sub-Gaussian distributions supported on $\{\sum_{j=1}^n \sigma_j v_j\}_{\sigma \in \{\pm 1\}^n}$, as established by [Dadush et al. \(2019\)](#). Their algorithm, known as the *Gram-Schmidt walk*, focuses on correlating the entries of σ in order to control the sub-Gaussian constant of $A\sigma$. Further, in applications of discrepancy theory such as randomized control trials, it is important to output not only a single signing but rather a *distribution* over low-discrepancy signings for the purpose of inferring treatment effects ([Krieger et al., 2019](#); [Turner et al., 2020](#); [Harshaw et al., 2021](#)).

To construct our relaxation, we replace the Rademacher distribution with the standard Gaussian distribution in the coupling interpretation (3) of combinatorial discrepancy. The *Gaussian discrepancy* is defined to be

$$\text{Gdisc}(A) = \min \left\{ \mathbb{E} \|Ag\|_\infty : g_j \sim \mathcal{N}(0, 1) \text{ for all } j \in [n], g_1, \dots, g_n \text{ jointly Gaussian} \right\}, \quad (4)$$

where $\mathcal{N}(0, 1)$ denotes the standard normal distribution. Equivalently, the minimization is over covariance matrices Σ for the random vector $g = (g_1, \dots, g_n)$ that lie in the *elliptope* \mathcal{E}_n , the set of all positive semidefinite matrices whose diagonal is the all-ones vector. Unlike Rademacher couplings which require, a priori, an exponential number of parameters to describe, joint Gaussian couplings admit a compact description that is completely determined by the covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$.

Let us record the simple observation that Gaussian discrepancy is indeed a relaxation of combinatorial discrepancy.

Proposition 3 *For any matrix A , it holds that $\text{Gdisc}(A) \leq \sqrt{2/\pi} \text{disc}(A)$.*

Proof Given an optimal signing σ for $\text{disc}(A)$, let $g_j = \sigma_j \xi$ for $j \in [n]$, where ξ is a standard Gaussian (equivalently, the covariance matrix of g is $\Sigma = \sigma \sigma^\top$). Then,

$$\text{Gdisc}(A) \leq \mathbb{E} \|Ag\|_\infty = \|A\sigma\|_\infty \mathbb{E}|\xi| = \sqrt{2/\pi} \text{disc}(A).$$

2. Such a minimization can be seen as an instance of a multimarginal optimal transport problem ([Pass, 2015](#); [Altschuler and Boix-Adserà, 2021](#)).

3. More precisely, in this statement we refer to the general result of [Banaszczyk \(1998\)](#) about balancing vectors to lie in a convex body K ; the case where K is the scaled ℓ_∞ ball corresponds to the combinatorial discrepancy.

■

As a consequence, results on ordinary discrepancy, such as Spencer's theorem and Banaszczyk's theorem, immediately translate into bounds on the Gaussian discrepancy.

1.2. Notation

We write $x \wedge y = \min\{x, y\}$, $x \vee y = \max\{x, y\}$, and $[n] = \{1, 2, \dots, n\}$. We write $\mathcal{N}(\mu, \Sigma)$ for the Gaussian distribution with mean μ and covariance matrix Σ . S^{n-1} denotes the unit sphere centered at the origin in \mathbb{R}^n , and \mathbb{S}_+^n denotes the set of symmetric positive semidefinite $n \times n$ matrices. For $A, B \in \mathbb{S}_+^n$ we write $A \preceq B$ if $B - A \in \mathbb{S}_+^n$. For vectors $x \in \mathbb{R}^n$ we write $\|x\|_p$, $p \in [1, \infty]$ for the ℓ_p norm of x , while for matrices X , $\|X\|_{p \rightarrow q}$ denotes the induced operator norm from ℓ_p to ℓ_q . We write $\langle \cdot, \cdot \rangle$ for both the Euclidean inner product between vectors and the Frobenius inner product between matrices. For two sequences of positive real numbers $(a_n), (b_n)$ we write $a_n \lesssim b_n$ if there exists a constant $C > 0$ with $a_n \leq Cb_n$ for all sufficiently large n , and write $a_n \asymp b_n$ if $a_n \lesssim b_n \lesssim a_n$. Thus, $a_n \lesssim b_n$ is a synonym for $a_n = O(b_n)$. The $n \times n$ identity matrix is denoted by I_n , and the all-ones vector is denoted by $\mathbf{1}_n \in \mathbb{R}^n$. We define the ellotope to be $\mathcal{E}_n = \{\Sigma \in \mathbb{R}^{n \times n} : \Sigma \succeq 0, \text{diag}(\Sigma) = \mathbf{1}_n\}$.

1.3. Results

1.3.1. COMPARISONS BETWEEN RELAXATIONS OF DISCREPANCY

In this section we develop an understanding of Gaussian discrepancy by comparing it to the *vector discrepancy* and *spherical discrepancy* relaxations of combinatorial discrepancy. Let us first define these notions and describe some results known about them.

The *vector discrepancy* relaxation replaces the signs $\sigma_1, \dots, \sigma_n \in \{\pm 1\}$ in combinatorial discrepancy (1) with unit vectors $u_1, \dots, u_n \in S^{n-1}$. Formally,

$$\text{Vdisc}(A) := \min \left\{ \max_{i \in [m]} \left\| \sum_{j=1}^n A_{i,j} u_j \right\|_2 : u_1, \dots, u_n \in S^{n-1} \right\}. \quad (5)$$

This problem is recast as a semidefinite program over the ellotope by constructing the Gram matrix, $\Sigma_{i,j} := \langle u_i, u_j \rangle$ for $i, j \in [n]$, of the unit vectors u_1, \dots, u_n ; see [Nikolov \(2013\)](#) for more details.

Given a signing $\sigma \in \{\pm 1\}^n$ and $u \in S^{n-1}$, we can associate to it the unit vectors $u_j = \sigma_j u$, $j \in [n]$. From this we see that vector discrepancy is indeed a relaxation of discrepancy:

$$\text{Vdisc}(A) \leq \text{disc}(A).$$

Vector discrepancy has been highly influential in discrepancy theory. It led to the initial algorithm of [Bansal \(2010\)](#) for Spencer's theorem, which uses a random walk guided by vector discrepancy solutions (see also [Bansal et al., 2016](#)), as well as the recent constructive proof of Banaszczyk's theorem ([Bansal et al., 2018](#)).⁴ Vector discrepancy has also been studied in its own right in the work of [Nikolov \(2013\)](#) which provides an in-depth analysis of this relaxation using SDP duality.

4. [Bansal et al. \(2018\)](#) remark that their Gram-Schmidt algorithm was "...inspired by the constructive proof of [Dadush et al. \(2019\)](#) for the existence of solutions to the Komlós vector coloring program of Nikolov."

Spherical discrepancy was more recently introduced by [Jones and McPartlon \(2020\)](#). It is obtained by relaxing the space of solutions in (1) from $\{\pm 1\}^n$ to the sphere of radius \sqrt{n} . Formally,

$$\text{Sdisc}(A) := \min_{x \in \sqrt{n} S^{n-1}} \|Ax\|_\infty. \quad (6)$$

[Jones and McPartlon \(2020\)](#) prove sharp bounds on spherical discrepancy in the setting of Spencer's theorem ($|A_{i,j}| \leq 1$ for all i, j) and the Komlós conjecture ($\|A_{:,j}\|_2 \leq 1$ for all $j \in [n]$) and apply their results to derive lower bounds for certain sphere covering problems.

With these definitions in hand, it is natural to wonder about the relationships between the various relaxations of discrepancy. For example, which relaxation gives the best approximation to combinatorial discrepancy? Our first result, whose proof we defer to Section 2, provides a unifying probabilistic perspective on the three relaxations of discrepancy and leads to a straightforward comparison.

Theorem 4 *For any matrix A , it holds that*

$$\text{Vdisc}(A) \asymp \min \left\{ \max_{i \in [m]} \mathbb{E}|(Ag)_i| \mid g \sim \mathcal{N}(0, \Sigma), \Sigma \in \mathbb{S}_+^n, \text{diag } \Sigma = \mathbf{1}_n \right\}, \quad (7)$$

$$\text{Sdisc}(A) \asymp \min \left\{ \mathbb{E} \max_{i \in [m]} |(Ag)_i| \mid g \sim \mathcal{N}(0, \Sigma), \Sigma \in \mathbb{S}_+^n, \text{tr } \Sigma = n \right\}. \quad (8)$$

Here, \mathbb{S}_+^n is the set of symmetric positive semidefinite $n \times n$ matrices.

For comparison, we recall that

$$\text{Gdisc}(A) = \min \left\{ \mathbb{E} \max_{i \in [m]} |(Ag)_i| \mid g \sim \mathcal{N}(0, \Sigma), \Sigma \in \mathbb{S}_+^n, \text{diag } \Sigma = \mathbf{1}_n \right\}.$$

Hence, Vdisc relaxes the objective of Gaussian discrepancy by moving the maximum over rows outside of the expectation, whereas Sdisc relaxes the constraint of Gaussian discrepancy from $\text{diag } \Sigma = \mathbf{1}_n$ to $\text{tr } \Sigma = n$. In comparison, from Proposition 3, the usual definition of discrepancy can be understood as *adding* a constraint to Gaussian discrepancy, namely that $\text{rank } \Sigma = 1$. The following relationship between the notions of discrepancy is an immediate consequence.

Corollary 5 *For any matrix A ,*

$$\text{Vdisc}(A) \vee \text{Sdisc}(A) \lesssim \text{Gdisc}(A) \lesssim \text{disc}(A).$$

In words, Gaussian discrepancy is a tighter relaxation of discrepancy than vector discrepancy and spherical discrepancy. Moreover, we show via a suite of examples in Section 2 that none of the inequalities in Corollary 5 can be reversed up to constant factor and that spherical discrepancy and vector discrepancy are incomparable in general.

Although Gaussian discrepancy is always larger than the vector discrepancy, the two relaxations are in fact closely related. Indeed, their common feasible set is the ellipope, which consists of the Gram matrices of collections of unit vectors or equivalently the covariance structures of feasible Gaussian couplings. The connection between these two relaxations yields our main tools for bounding Gaussian discrepancy as well as an $O(\sqrt{\log m})$ -factor approximation algorithm for computing Gaussian discrepancy.

Our main inequality for controlling Gaussian discrepancy relies on a notion of *rank-constrained vector discrepancy*, defined as follows. For every rank $r \in [n]$, let

$$\text{Vdisc}_r(A) := \min \left\{ \max_{i \in [m]} \left\| \sum_{j=1}^n A_{i,j} u_j \right\|_2 : u_1, \dots, u_n \in S^{r-1} \right\}. \quad (9)$$

Equivalently we require the Gram matrix $\Sigma := (\langle u_i, u_j \rangle)_{i,j \in [n]}$ to have rank at most r . Since a rank-1 matrix Σ corresponds precisely to a signing, note that

$$\text{Vdisc}(A) = \text{Vdisc}_n(A) \leq \text{Vdisc}_{n-1}(A) \leq \dots \leq \text{Vdisc}_1(A) = \text{disc}(A).$$

The next result compares Gaussian discrepancy with this rank-constrained problem and is a crucial ingredient in our study of online Gaussian discrepancy.

Proposition 6 *For any $r \in [n]$, it holds that $\text{Gdisc}(A) \leq \sqrt{r} \text{Vdisc}_r(A)$.*

Proof Let $u_1, \dots, u_n \in S^{r-1}$ be an optimal solution for $\text{Vdisc}_r(A)$. Then, if ξ denotes a standard Gaussian vector in \mathbb{R}^r , we can define $g_j := \langle u_j, \xi \rangle$ for $j \in [n]$, and it can be easily checked that this is feasible for the definition of Gaussian discrepancy. Hence,

$$\text{Gdisc}(A) \leq \mathbb{E} \max_{i \in [m]} \left| \left\langle \sum_{j=1}^n A_{i,j} u_j, \xi \right\rangle \right| \leq \max_{i \in [m]} \left\| \sum_{j=1}^n A_{i,j} u_j \right\|_2 \mathbb{E} \|\xi\|_2 \leq \sqrt{r} \text{Vdisc}_r(A). \quad (10)$$

The last inequality uses the standard probabilistic fact that $\mathbb{E} \|\xi\|_2 \leq \sqrt{r}$. When $r = 1$, this can be improved to $\mathbb{E} |\xi| = \sqrt{2/\pi}$, which recovers Proposition 3. \blacksquare

Applying a similar strategy but using the union bound instead of Cauchy–Schwarz shows that vector discrepancy upper bounds Gaussian discrepancy up to a logarithmic term.

Proposition 7 *It holds that $\text{Gdisc}(A) \lesssim \sqrt{\log(2m)} \text{Vdisc}(A)$.*

Proof Let $u_1, \dots, u_n \in \mathbb{R}^n$ be an optimal solution for $\text{Vdisc}_n(A)$, and let $\xi \sim \mathcal{N}(0, I_n)$. As in the proof of Proposition 6,

$$\text{Gdisc}(A) \leq \mathbb{E} \max_{i \in [m]} \left| \left\langle \sum_{j=1}^n A_{i,j} u_j, \xi \right\rangle \right|.$$

The random variable $\langle \sum_{j=1}^n A_{i,j} u_j, \xi \rangle$ is sub-Gaussian with parameter $\|\sum_{j=1}^n A_{i,j} u_j\|_2^2$. Thus the standard maximal inequality for sub-Gaussian random variables (see, e.g., Boucheron et al., 2013, Theorem 2.5) yields

$$\text{Gdisc}(A) \leq \sqrt{2 \ln(2m)} \max_{i \in [m]} \left\| \sum_{j=1}^n A_{i,j} u_j \right\|_2 = \sqrt{2 \ln(2m)} \text{Vdisc}_n(A). \quad (11)$$

\blacksquare

We record three useful consequences of Proposition 7. When combined with Theorem 4 it implies that Gaussian discrepancy is approximated by vector discrepancy up to a logarithmic factor. Since

vector discrepancy admits an SDP formulation, it can be computed in polynomial time using methods from convex optimization (Nesterov, 2018), and in turn this yields our claimed $O(\sqrt{\log m})$ -factor approximation algorithm for Gaussian discrepancy.

Second, Proposition 7 allows us to draw a new connection between spherical discrepancy and vector discrepancy. Combining this result with Corollary 5 and Proposition 6, we obtain the following.

Corollary 8 *For any matrix A , it holds that $\text{Sdisc}(A) \lesssim \sqrt{n \wedge \log m} \text{Vdisc}(A)$.*

As described in detail in Section 2, this inequality is tight and also may not be reversed — in fact it is possible to have $\text{Sdisc}(A) = 0$ while $\text{Vdisc}(A) > 0$ (see Example 2). We remark that this sharp result falls out naturally with Gaussian discrepancy as a mediator between spherical and vector discrepancy.

Finally, Proposition 7 leads to a simple algorithm that achieves Banaszczyk’s bound for Gaussian discrepancy. The main result of Nikolov (2013) states that the Komlós conjecture holds for vector discrepancy. Hence, a solution to the SDP formulation of vector discrepancy yields a feasible covariance matrix for Gaussian discrepancy with objective value $O(\sqrt{\log m})$.⁵ Note that since Gaussian discrepancy is a relaxation, existing approaches for combinatorial discrepancy also achieve Banaszczyk’s bound for Gaussian discrepancy with much more involved algorithms (Bansal et al., 2016; Levy et al., 2017; Bansal et al., 2018).

This raises the questions of whether or not the Komlós conjecture (Conjecture 1) or online Banaszczyk (Conjecture 2) can be proved for Gaussian discrepancy. More broadly, it is natural to ask whether unresolved conjectures about combinatorial discrepancy can be solved for a given relaxation. Indeed, Nikolov (2013) and Jones and McPartlon (2020) establish the Komlós conjecture for vector and spherical discrepancy, respectively. While we are not able to establish a Gaussian version of the Komlós conjecture, it serves as a tantalizing question for further study; see Section 1.4 for further discussion. On the other hand, our main algorithmic result establishes a Gaussian discrepancy variant of the online Banaszczyk bound, as discussed in the next section.

1.3.2. ONLINE BANASZCZYK BOUND FOR GAUSSIAN DISCREPANCY

We begin by recalling the setting of online discrepancy minimization with an oblivious adversary. First, an adversary picks vectors v_1, \dots, v_T of ℓ_2 norm at most 1. Then, in each round $t = 1, 2, \dots, T$ the vector v_t is revealed to the algorithm, which must output a sign $\sigma_t \in \{\pm 1\}$. The aim of the algorithm is to minimize the maximum discrepancy incurred, i.e. $\max_{t \in [T]} \|\sum_{s=1}^t \sigma_s v_s\|_\infty$. Equivalently, in each round the algorithm is required to choose a *vector* of signs $\sigma^{(t)} \in \{\pm 1\}^t$ that obeys the following *consistency condition*:

$$\sigma_{[t-1]}^{(t)} = \sigma^{(t-1)}, \quad (12)$$

where $x_{[t-1]}$ denotes the restriction of $x \in \mathbb{R}^t$ to the first $t-1$ coordinates. That is, the new signing $\sigma^{(t)} \in \{\pm 1\}^t$ chosen in round t is forbidden from changing the signs specified in previous rounds.

The latter formulation motivates our Gaussian discrepancy variant of the online discrepancy problem. At each round t , we require the algorithm to output a t -dimensional correlation matrix $\Sigma^{(t)}$ which satisfies a consistency condition analogous to (12), namely

$$\Sigma_{[t-1] \times [t-1]}^{(t)} = \Sigma^{(t-1)}. \quad (13)$$

5. This can be further refined to $O(\sqrt{\log(m \wedge n)})$ using standard reductions.

Thus after round t , the correlations between the first t coordinates must remain fixed. The aim is still to minimize the maximum discrepancy incurred, i.e. $\max_{t \in [T]} \mathbb{E} \|\sum_{s=1}^t g_s v_s\|_\infty$, where $g = (g_1, \dots, g_T) \sim \mathcal{N}(0, \Sigma^{(T)})$. A simple argument based on the Cholesky decomposition (see Appendix A for details) shows that any such sequence $\Sigma^{(1)}, \dots, \Sigma^{(T)}$ can be realized as a sequence of Gram matrices of initial segments of a sequence of unit vectors $u_1, \dots, u_T \in \ell_2$. Hence, it is equivalent to require that the algorithm output a unit vector $u_t \in \ell_2$ at each round t , and then set $\Sigma_{i,j}^{(t)} = \langle u_i, u_j \rangle_{\ell_2}$. We arrive at the following formulation of online Gaussian discrepancy. As in online combinatorial discrepancy, we allow the algorithm to be randomized.

Problem 1 (Online Gaussian discrepancy with oblivious adversary) *An adversary selects vectors v_1, \dots, v_T with ℓ_2 norm at most 1 in advance. At each round $t \in [T]$, the algorithm observes v_t and outputs a random unit vector $u_t \in \ell_2$. The algorithm aims to minimize*

$$\text{Gdisc}(v_1, \dots, v_T; \Sigma^{(T)}) = \max_{t \in [T]} \mathbb{E} \left\| \sum_{s=1}^t g_s v_s \right\|_\infty, \quad g \sim \mathcal{N}(0, \Sigma^{(T)})$$

with high probability over the random correlation matrix $\Sigma_{i,j}^{(T)} = \langle u_i, u_j \rangle_{\ell_2}$.

We show in Appendix A that, without loss of generality, it suffices to consider u_t whose support is a subset of the first t coordinates.

One strategy for generating feasible couplings in each round is to first fix a *rank* parameter $r \geq 1$ in advance and output a unit vector $u_t \in S^{r-1}$ in each round. Our main result is an algorithm of this form that solves the Gaussian discrepancy variant of the online Banaszczyk conjecture (Conjecture 2).

Theorem 9 (Online Banaszczyk bound for Gaussian discrepancy) *Let $v_1, \dots, v_T \in \mathbb{R}^m$ denote vectors of ℓ_2 norm at most 1 selected in advance by an adversary. For all positive integers $r \geq 2$, there is a randomized online algorithm that with probability at least $1 - \delta$ outputs $u_1, \dots, u_T \in S^{r-1}$ such that*

$$\text{Gdisc}(v_1, \dots, v_T; \Sigma^{(T)}) = O(\sqrt{\log(mT/\delta)}),$$

where $\Sigma_{i,j}^{(T)} = \langle u_i, u_j \rangle$. The algorithm runs in time $O(mr)$ per round.

The proof of Theorem 9 is the main content of Section 3. If we could prove this theorem with $r = 1$, then the online Banaszczyk problem (Conjecture 2) would follow from the proof of Proposition 3. Unfortunately, there is an obstruction preventing us from considering $r = 1$, as described further below. Nevertheless, Theorem 9 shows that we are ‘one rank away’ from establishing Conjecture 2.

Our algorithm is based on an intriguing idea in the recent paper Liu et al. (2022): namely, if one can find a Markov chain on \mathbb{R} whose increments take values in $\{\pm 1\}$ and whose stationary distribution is a Gaussian, then it is possible to construct an algorithm, which they call the *Gaussian fixed-point walk*, with the property that each partial sum $\sum_{s=1}^t \sigma_s v_s$ is the difference of two Gaussian vectors. The online Banaszczyk bound would then follow from a union bound. However, Liu et al. (2022) exhibit a *parity obstruction* which implies no such Markov chain exists on \mathbb{R} , and this leads them to consider instead Markov chains whose increments lie in $\{0, \pm 1\}$ or $\{\pm 1, 2\}$ (and hence the resulting algorithms output partial colorings or improper colorings). We show that an analogous Markov chain *does* exist on \mathbb{R}^r for any $r \geq 2$ whose increments lie in S^{r-1} and whose stationary

distribution is a Gaussian with independent and identically distributed coordinates (see Figure 1). Working with higher ranks $r \geq 2$ avoids certain technical complications resulting in the partial coloring version derived by Liu et al. (2022) and allows for a simple algorithm and analysis.

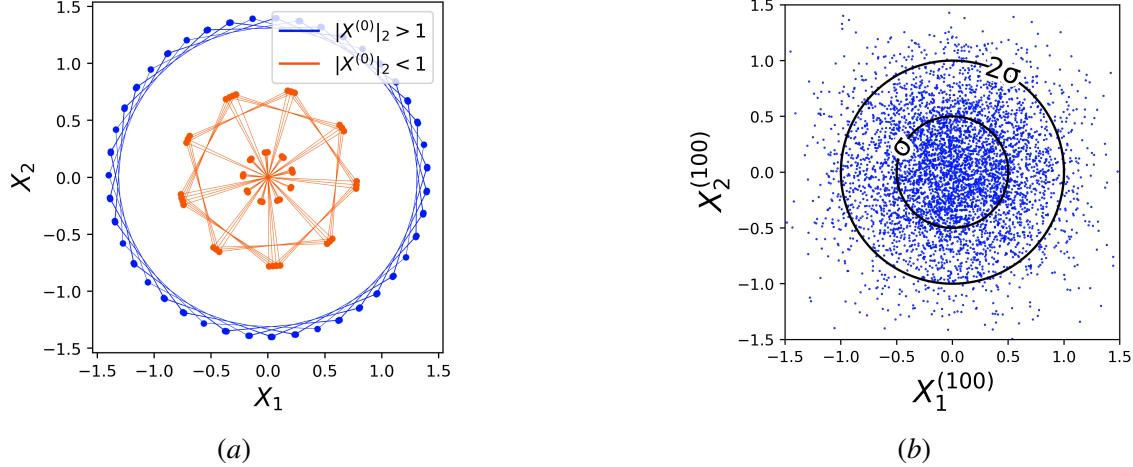


Figure 1: $(X^{(t)})_{t \geq 0}$ is the Markov chain on \mathbb{R}^2 with unit vector steps and Gaussian stationary distribution. (a) Evolution of $(X^{(t)})_{t \geq 0}$ over two runs of the Markov chain with 2000 steps. The orange and blue trajectories are initialized inside and outside the unit circle, respectively. (b) Scatter plot of $X^{(100)}$ over 5000 independent runs of the Markov chain started from the stationary distribution $X^{(0)} \sim \mathcal{N}(0, \sigma^2 I_2)$, where $\sigma = 0.5$.

We note that the guarantee of Theorem 9 can be shown to be tight by using standard Gaussian estimates. In particular, taking $r > 2$ in our algorithm provably does not improve the Gaussian discrepancy.

Interestingly, our algorithm also has consequences for the previously unexplored problem of online vector discrepancy, and for this problem, choosing $r > 2$ offers a substantial improvement. Perhaps surprisingly, Theorem 10 below shows that the Komlós conjecture, with sharp constant, is attainable for vector discrepancy even in the oblivious online setting. In contrast, Bansal et al. (2020) exhibit an $\tilde{\Omega}(\sqrt{\log T})$ lower bound for online combinatorial discrepancy in the oblivious setting. The next result is an immediate consequence of Theorem 15 proved in Section 3.

Theorem 10 (Online Komlós bound for vector discrepancy) *Let $v_1, \dots, v_T \in \mathbb{R}^m$ denote vectors of ℓ_2 norm at most 1 selected in advance by an adversary. Fix $\delta, \varepsilon > 0$. When run with rank $r = \Theta(\varepsilon^{-2} \log(mT/\delta))$, the algorithm outputs $u_1, \dots, u_T \in S^{r-1}$ online with*

$$\max_{t \in [T]} \max_{i \in [m]} \left\| \sum_{s=1}^t (v_s)_i u_s \right\|_2 \leq 1 + \varepsilon$$

with probability at least $1 - \delta$. The algorithm runs in time $O(m\varepsilon^{-2} \log(mT/\delta))$ per round.

As a corollary of the above theorem, we obtain a new proof of the vector Komlós theorem of Nikolov (2013), which states that $\text{Vdisc}(A) \leq 1$ for all matrices A whose columns have ℓ_2 norm at most 1. Since the vector discrepancy of the identity matrix I_T is 1, our result is essentially sharp.

Although we present our algorithm for the online setting, it yields new algorithmic implications for the offline setting as well. In particular, our runtime is nearly linear in the input size mT for moderate values of the approximation parameter ε . This improves on the time complexity of previous algorithms for vector Komlós such as off-the-shelf SDP solvers, which require $\tilde{O}((m \vee T)^{3.5} \log(1/\varepsilon))$ arithmetic operations (Nesterov, 2018), and an iterative approach of Dadush et al. (2019) that naïvely runs in time $O(T^3(m \vee T))$ due to costly matrix inversion steps.⁶

Furthermore, setting $\varepsilon = 1$ in Theorem 10 implies that the Komlós conjecture holds for rank-constrained vector discrepancy with $r = O(\log(mT))$. To the best of our knowledge this result was previously unknown even in an offline setting.

1.4. Open problems

Our work leads to some natural open questions which we briefly describe below.

1. **Komlós conjecture for Gaussian discrepancy.** Does there exist a constant $K_G > 0$ such that $Gdisc(A) \leq K_G$ for any matrix A whose columns have ℓ_2 norm at most 1?

Since Gaussian discrepancy is a relaxation, solving this conjecture is a natural prerequisite for establishing the original Komlós conjecture (Conjecture 1). It also leads to a related problem: prove that the Komlós conjecture for Gaussian discrepancy implies the Komlós conjecture.

2. **Rounding Gaussian discrepancy solutions.** Does there exist an efficient rounding scheme to convert Gaussian discrepancy solutions to low-discrepancy signings? For example, in the Komlós setting, does there exist a polynomial-time computable function f from covariance matrices to signings, such that

$$\|Af(\Sigma)\|_\infty \lesssim \mathbb{E}_{g \sim \mathcal{N}(0, \Sigma)} \|Ag\|_\infty + O(1) ,$$

for all matrices A with $\max_{j \in [n]} \|A_{:,j}\|_2 \leq 1$?

In Appendix B, we show that two simple rounding schemes (Goemans–Williamson rounding and a PCA-based rounding) are not effective in the setting of Spencer’s theorem and the Komlós conjecture.

3. **Computational tractability of Gaussian discrepancy.** Given a matrix, can we compute its Gaussian discrepancy exactly or approximately in polynomial time? In particular, is it NP-hard to approximate Gaussian discrepancy up to a constant factor? We note that hardness of approximation results are already known for combinatorial discrepancy (Charikar et al., 2011) and spherical discrepancy (Jones and McPartlon, 2020).
4. **Achieving Banaszczyk’s bound online.** Finally, we mention that the question of achieving Banaszczyk’s bound online for combinatorial discrepancy, as originally posed by Alweiss et al. (2021), is still open.

6. We do not take into account fast matrix multiplication for these runtime estimates. We also note that an SDP solver would in general yield the stronger guarantee $\|\sum_{s=1}^T (v_t)_i u_t\|_2^2 \leq Vdisc(A)^2 + \varepsilon$ for $1 \leq i \leq m$.

1.5. Organization of the paper

In Section 2.1, we give the proof of our main comparison result between the relaxations of discrepancy (Theorem 4). In Section 2.2, we provide a suite of examples showing that the inequalities given in Section 1.3.1 are sharp. Finally, in Section 3, we present our algorithm and prove that it achieves the online Banaszczyk bound for Gaussian discrepancy.

2. COMPARISONS BETWEEN RELAXATIONS OF DISCREPANCY

In this section we prove Theorem 4 and give a number of illustrative examples that highlight the differences between the various relaxations of discrepancy. These examples show that none of the inequalities in Corollary 5 can be reversed in general (at least with constants independent of m and n). Moreover, our results in this section give evidence for the following assertion: vector discrepancy and spherical discrepancy each capture distinct aspects of the original discrepancy problem, namely SDP solutions are “aligned with the coordinate axes”, whereas spherical discrepancy solutions are “low rank”. Gaussian discrepancy appears to capture both of these aspects simultaneously.

This assertion can already be partially justified by observing that the feasible covariance matrices Σ for vector discrepancy satisfy $\text{diag } \Sigma = \mathbf{1}_n$, i.e. the variance along each of the coordinate axes is fixed to be 1. On the other hand, from the proof of Theorem 4 below, we see that a spherical discrepancy solution x gives rise to the covariance matrix $\Sigma := xx^\top$, which is indeed low rank but is not guaranteed to have unit variance along each coordinate axis.

2.1. Proof of main comparison result

We repeatedly use the following useful facts about log-concave distributions. Note that any Gaussian random variable has a log-concave distribution.

Lemma 11 (Alonso-Gutiérrez and Bastero (2015, Propositions A.5-A.6)) *Let X be a random vector in \mathbb{R}^d with a log-concave distribution, and let $\|\cdot\| : \mathbb{R}^d \rightarrow \mathbb{R}_+$ be any seminorm.*

1. *For any $p > 1$, it holds that $\mathbb{E}[\|X\|^p]^{1/p} \lesssim p \mathbb{E}\|X\|$, where the implied constant is universal.*
2. *For any $t > 0$, it holds that $\mathbb{P}\{\|X\| \leq t \mathbb{E}\|X\|\} \lesssim t$, where the implied constant is universal.*

Proof [Proof of Theorem 4] Vector discrepancy: Given $\Sigma \in \mathcal{E}_n$, let g^Σ denote a centered Gaussian vector with covariance matrix Σ , and write $UU^\top = \Sigma$ with u_j denoting the rows of U . Observe that

$$\max_{i \in [m]} \left\| \sum_{j=1}^n A_{i,j} u_j \right\|_2 = \sqrt{\max_{i \in [m]} \langle A_{i,:}, \Sigma A_{i,:} \rangle}. \quad (14)$$

Therefore

$$\begin{aligned} \text{Vdisc}(A) &= \min_{\Sigma \in \mathcal{E}_n} \sqrt{\max_{i \in [m]} \langle A_{i,:}, \Sigma A_{i,:} \rangle} = \min_{\Sigma \in \mathcal{E}_n} \sqrt{\max_{i \in [m]} \mathbb{E} \left[\left| \sum_{j=1}^n A_{i,j} g_j^\Sigma \right|^2 \right]} \\ &\stackrel{(*)}{\asymp} \min_{\Sigma \in \mathcal{E}_n} \max_{i \in [m]} \mathbb{E} \left| \sum_{j=1}^n A_{i,j} g_j^\Sigma \right| = \min_{\Sigma \in \mathcal{E}_n} \max_{i \in [m]} \mathbb{E} |(Ag^\Sigma)_i|, \end{aligned}$$

where (\star) uses Lemma 11 and Cauchy–Schwarz.

Spherical discrepancy: Let us temporarily denote

$$\widetilde{\text{disc}}_s(A) := \min\{\mathbb{E}\|Ag\|_\infty : g \sim \mathcal{N}(0, \Sigma), \Sigma \in \mathbb{S}_+^n, \text{tr } \Sigma = n\}.$$

Let g be a Gaussian that achieves the minimum in the definition of $\widetilde{\text{disc}}_s(A)$. Then, by Markov's inequality,

$$\mathbb{P}\{\|Ag\|_\infty \geq 3\widetilde{\text{disc}}_s(A)\} \leq \frac{1}{3}. \quad (15)$$

From Lemma 11, we further have

$$\mathbb{P}\{\|g\|_2 \leq c\mathbb{E}\|g\|_2\} \leq \frac{1}{3} \quad \text{and} \quad \mathbb{E}\|g\|_2 \gtrsim \sqrt{\mathbb{E}\|g\|_2^2} \quad (16)$$

for some constant $c > 0$. From (16) and (15), there exists a realization $g \in \mathbb{R}^n$ of g with

$$\|Ag\|_\infty \leq 3\widetilde{\text{disc}}_s(A), \quad \text{and} \quad \|g\|_2 \geq c\mathbb{E}\|g\|_2 \gtrsim \sqrt{\mathbb{E}\|g\|_2^2} = \sqrt{n}.$$

It follows that

$$\text{Sdisc}(A) \leq \left\| A \frac{\sqrt{n}g}{\|g\|_2} \right\|_\infty = \frac{\sqrt{n}}{\|g\|_2} \|Ag\|_\infty \lesssim \widetilde{\text{disc}}_s(A).$$

Conversely, let x be an optimal solution for $\text{Sdisc}(A)$ and let ξ be a standard Gaussian variable on \mathbb{R} , so $\xi x \sim \mathcal{N}(0, xx^\top)$ and $\text{tr}(xx^\top) = \|x\|_2^2 = n$. Then,

$$\widetilde{\text{disc}}_s(A) \leq \mathbb{E}\|A(\xi x)\|_\infty = \mathbb{E}|\xi| \|Ax\|_\infty = \sqrt{\frac{2}{\pi}} \text{Sdisc}(A),$$

which establishes the converse bound. ■

2.2. Tightness of the comparison results

Our first example demonstrates the sharpness of Corollary 8 with respect to both m and n .

Example 1 For this example we first consider an infinitely tall matrix A . Let the rows of A consist of all unit vectors in \mathbb{R}^n . By taking the vector discrepancy solution $\Sigma = I_n$ and using (14), we see that

$$\text{Vdisc}(A) \leq \sqrt{\max_{a \in S^{n-1}} \langle a, \Sigma a \rangle} = 1.$$

(In fact it is not hard to see that this is an equality.) However, for any $x \in \sqrt{n}S^{n-1}$ we have

$$\max_{a \in S^{n-1}} |\langle a, x \rangle| = \|x\|_2 = \sqrt{n},$$

which shows that $\text{Sdisc}(A) = \sqrt{n}$. Thus, the spherical discrepancy can be much larger than vector discrepancy.

Although this example used an infinitely tall matrix A , we can modify it by taking the rows of A to consist of a $1/2$ -net of S^{n-1} . Then, the number of rows of A can be taken to be $m = \exp(\Theta(n))$, and a standard argument involving nets (see the proof of [Vershynin, 2018, Lemma 4.4.1](#)) shows that we still have $\text{Vdisc}(A) \leq 1$ and $\text{Sdisc}(A) \gtrsim \sqrt{n}$. In particular, this shows that the bound

$$\text{Sdisc}(A) = O(\sqrt{n \wedge \log m}) \text{Vdisc}(A)$$

obtained in [Corollary 8](#) is sharp with respect to both m and n .

The next example shows that [Corollary 8](#) cannot be reversed, even with a constant depending on m and n , so that vector discrepancy cannot in general be controlled by spherical discrepancy.

Example 2 Let $v \in S^{n-1}$ be a unit vector and suppose that the rows of A consist of all vectors in

$$\mathcal{A} := S^{n-1} \cap v^\perp,$$

the set of unit vectors which are orthogonal to v (as in the preceding example, this example can also be modified to a matrix with finitely many rows). Then, $\text{Sdisc}(A) \leq \sqrt{n} \max_{a \in \mathcal{A}} |\langle a, v \rangle| = 0$. On the other hand, we claim that $\text{Vdisc}(A) > 0$ for most choices of v . Suppose to the contrary that there exist unit vectors $u_1, \dots, u_n \in \mathbb{R}^n$ witnessing the fact that $\text{Vdisc}(A) = 0$. Then, for each $k \in [n]$,

$$0 = \max_{a \in \mathcal{A}} \left\| \sum_{j=1}^n a_j u_j \right\|_2 \geq \max_{a \in \mathcal{A}} \left| \sum_{j=1}^n a_j u_j[k] \right| = \max_{a \in \mathcal{A}} |\langle a, u_\bullet[k] \rangle| = \|\text{proj}_{v^\perp} u_\bullet[k]\|_2,$$

where $u_\bullet[k] \in \mathbb{R}^n$ denotes the vector $(u_j[k])_{j=1}^n$. The inequality implies that $u_\bullet[k]$ is a multiple of v , i.e., there exists a scalar $c_k \in \mathbb{R}$ such that $u_j[k] = c_k v_j$ for all $j \in [n]$. Since u_j is a unit vector,

$$1 = \|u_j\|_2^2 = v_j^2 \sum_{k=1}^n c_k^2, \quad \text{for all } j \in [n].$$

In order for this to hold, each coordinate of v must have the same magnitude, i.e. v must be a signing (scaled by $n^{-1/2}$). Hence, for most directions v we have $\text{Sdisc}(A) = 0$ but $\text{Vdisc}(A) > 0$.

This shows in particular that there does not exist $C > 0$ such that $\text{Vdisc}(A) \leq C \text{Sdisc}(A)$ for all $m \times n$ matrices A , even if the constant C is allowed to depend on m and n .

The two preceding examples show that the statements $\text{Vdisc} \leq C_1 \text{Sdisc}$ and $\text{Sdisc} \leq C_2 \text{Vdisc}$ do not hold with universal constants $C_1, C_2 > 0$. In particular, since $\text{Vdisc} \vee \text{Sdisc} \lesssim \text{Gdisc}$ by [Corollary 5](#), it also implies that the statements $\text{Gdisc} \leq C'_1 \text{Vdisc}$ and $\text{Gdisc} \leq C'_2 \text{Sdisc}$ also do not hold with universal constants C'_1 and C'_2 .

We give another example to show that in general, Gaussian discrepancy can indeed be smaller than combinatorial discrepancy.

Example 3 Consider the case when $m = 1$, so that A consists of a single row, and assume that the entries of A are i.i.d. standard Gaussians. Then, the discrepancy of $A = (a_1, \dots, a_n)$ is usually referred to as the number balancing problem.

Let us suppose that n is a multiple of 3, and divide the n coordinates into 3 groups consisting of $n/3$ coordinates each. Due to concentration of measure, the sums

$$\ell_1 := \sum_{i=1}^{n/3} |a_i|, \quad \ell_2 := \sum_{i=n/3+1}^{2n/3} |a_i|, \quad \text{and} \quad \ell_3 := \sum_{i=2n/3+1}^n |a_i|$$

will concentrate around their expected value $\sqrt{2/\pi} n/3$. In particular, with high probability we will have $2 \max_{i=1,2,3} \ell_i \leq \sum_{i=1}^3 \ell_i$, which says that ℓ_1, ℓ_2 , and ℓ_3 form the side lengths of a triangle in \mathbb{R}^2 . If we let $u_1, u_2, u_3 \in \mathbb{R}^2$ denote unit vectors corresponding to the sides of these triangles, then

$$\begin{aligned} 0 &= \sum_{i=1}^{n/3} |a_i| u_1 + \sum_{i=n/3+1}^{2n/3} |a_i| u_2 + \sum_{i=2n/3+1}^n |a_i| u_3 \\ &= \sum_{i=1}^{n/3} a_i (\operatorname{sgn} a_i) u_1 + \sum_{i=n/3+1}^{2n/3} a_i (\operatorname{sgn} a_i) u_2 + \sum_{i=2n/3+1}^n a_i (\operatorname{sgn} a_i) u_3. \end{aligned}$$

This shows that with high probability, $\operatorname{Vdisc}_2(A) = 0$ and hence $\operatorname{Gdisc}(A) = 0$ by Proposition 6.

On the other hand, it is well-known that $\operatorname{disc}(A) = \Theta(\sqrt{n} 2^{-n}) > 0$ (see, e.g., [Karmarkar et al., 1986](#); [Costello, 2009](#); [Turner et al., 2020](#)). In particular, there does not exist a constant $C > 0$ (even if the constant is allowed to depend on m and n) such that $\operatorname{disc}(A) \leq C \operatorname{Gdisc}(A)$.

In summary, if we require the constants to be universal, then in general none of the inequalities in Corollary 5 can be reversed, and furthermore vector discrepancy and spherical discrepancy are incomparable.

In this section, we have argued that vector discrepancy and spherical discrepancy capture distinct aspects of the original discrepancy problem, and can therefore be viewed as complementary. It is then natural to ask, whether a *combination* of vector discrepancy and spherical discrepancy can control Gaussian discrepancy. This is also motivated by the results of [Nikolov \(2013\)](#) and [Jones and McPartlon \(2020\)](#) that prove the Komlós conjecture for vector discrepancy and spherical discrepancy respectively; hence, a control on the Gaussian discrepancy in terms of these two notions would imply the Komlós conjecture for Gaussian discrepancy as well. We however resolve this question negatively via the following example.

Example 4 Let the rows of A consist of all unit vectors in \mathbb{R}^n orthogonal to e_1 . Then, similarly to the previous examples, we have $\operatorname{Vdisc}(A) \leq 1$ and $\operatorname{Sdisc}(A) = 0$. On the other hand, for any feasible Gaussian coupling g ,

$$\mathbb{E} \|Ag\|_\infty = \mathbb{E} \max_{a \in S^{n-1} \cap e_1^\perp} |\langle a, g \rangle| = \mathbb{E} \|\operatorname{proj}_{e_1^\perp} g\|_2 \stackrel{(*)}{\gtrsim} \sqrt{\mathbb{E}[\|(0, g_2, \dots, g_n)\|_2^2]} = \sqrt{n-1},$$

where the inequality $(*)$ uses Lemma 11. Hence, $\operatorname{Gdisc}(A) \gtrsim \sqrt{n-1}$. This in fact shows that there is no non-trivial function $f : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ which is independent of both m and n , such that

$$\operatorname{Gdisc}(A) \leq f(\operatorname{Vdisc}(A), \operatorname{Sdisc}(A))$$

for all matrices A .

3. GAUSSIAN FIXED-POINT WALK IN HIGHER DIMENSIONS

3.1. High-level overview

We begin by summarizing the Gaussian fixed-point walk as introduced in [Liu et al. \(2022\)](#). Recall the setting of online discrepancy minimization with oblivious adversary: the adversary chooses in advance vectors $v_1, \dots, v_T \in \mathbb{R}^m$ and at each time step $t \in [T]$, the vector v_t is revealed to the algorithm, upon which it must then choose a sign $\sigma_t \in \{\pm 1\}$. The loss incurred by the algorithm is the maximum discrepancy $\max_{t \in [T]} \|\sum_{s=1}^t \sigma_s v_s\|_\infty$.

Let $w_t := w_0 + \sum_{s=1}^t \sigma_s v_s$ denote the partial sum vector with initialization $w_0 \sim \mathcal{N}(0, I_m)$. The idea behind the Gaussian fixed-point walk is to choose the signs σ_t to ensure that $w_t \sim \mathcal{N}(0, I_m)$ for all $t \in [T]$ (observe that if this holds, then the loss $\max_{t \in [T]} \|w_t - w_0\|_\infty$ incurred by the algorithm is easily controlled via a union bound). Towards this end, we may assume as an inductive hypothesis that $w_{t-1} \sim \mathcal{N}(0, I_m)$. Then, we may write

$$w_t = \underbrace{\text{proj}_{v_t^\perp} w_{t-1}}_{\sim \mathcal{N}(0, I_m - v_t v_t^\perp)} + \underbrace{[\text{proj}_{v_t} w_{t-1} + \sigma_t v_t]}_{\sim \mathcal{N}(0, v_t v_t^\perp)}.$$

If we choose σ_t independently of $\text{proj}_{v_t^\perp} w_{t-1}$, then the two terms in the above decomposition are independent. Moreover, if we could choose σ_t such that $\text{proj}_{v_t} w_{t-1} + \sigma_t v_t \sim \mathcal{N}(0, v_t v_t^\perp)$, then $w_t \sim \mathcal{N}(0, I_m)$. These considerations lead to the problem of finding a one-dimensional Markov chain (on $\text{span } v_t$) whose increments belong to $\{\pm v_t\}$, and whose stationary distribution is $\mathcal{N}(0, v_t v_t^\perp)$. However, a parity argument given in [Liu et al. \(2022\)](#) shows that such a Markov chain does not exist, even if we allow for changing the variance of the Gaussian from 1.

To deal with this issue, [Liu et al. \(2022\)](#) show that such Markov chains exist if we allow the increments to take values in $\{0, \pm v_t\}$ or $\{\pm v_t, 2v_t\}$, which leads to the algorithm outputting partial or improper colorings. If the variance of the Gaussian is set to be sufficiently large, then they can further argue that the algorithm outputs an actual signing with high probability, but the large variance of the Gaussian prevents them from establishing an online Banaszczyk result.

Our contribution in this section lies in adapting the idea above to online Gaussian discrepancy, by constructing a Markov chain in \mathbb{R}^r , $r \geq 2$ whose stationary distribution is a Gaussian, and whose increments lie in on the unit sphere; this is given in Section 3.2. The algorithm and analysis are then given in Section 3.3.

3.2. Construction of the Markov chain

Our goal in this section is to construct a Markov chain on \mathbb{R}^r for $r \geq 2$ whose increments are unit vectors and whose stationary distribution is $\mathcal{N}(0, \sigma^2 I_r)$ for some σ^2 . In this and subsequent sections, σ is used to denote the standard deviation (and in particular should not be confused with a signing).

The density of the χ -distribution with r degrees of freedom is

$$\chi_r(s) = \frac{1}{2^{r/2-1} \Gamma(r/2)} s^{r-1} e^{-\frac{s^2}{2}} \mathbb{1}\{s \geq 0\}.$$

Hence, the density of $\|g\|_2$ where $g \sim \mathcal{N}(0, \sigma^2 I_r)$ is

$$\chi_{r, \sigma^2}(s) = \frac{1}{2^{r/2-1} \Gamma(r/2)} \frac{s^{r-1}}{\sigma^r} e^{-\frac{s^2}{2\sigma^2}} \mathbb{1}\{s \geq 0\}.$$

Let $x \in \mathbb{R}^r$, and define the following sets:

$$S_x := \{y \in \mathbb{R}^r : \|x - y\|_2 = 1 \text{ and } \|y\|_2 = \|x\|_2\},$$

$$S'_x := \{y \in \mathbb{R}^r : \|x - y\|_2 = 1 \text{ and } \|y\|_2 = 1 - \|x\|_2\}.$$

Since $\|y\|_2^2 = \|x\|_2^2 + \|x - y\|_2^2 + 2\langle y - x, x \rangle$, if $\|y - x\|_2 = 1$ then from the Cauchy-Schwarz inequality, we obtain

$$(\|x\|_2 - 1)^2 \leq \|y\|_2^2 \leq (\|x\|_2 + 1)^2.$$

Conversely, if $s \geq 0$ satisfies $(\|x\|_2 - 1)^2 \leq s^2 \leq (\|x\|_2 + 1)^2$, then there exists $y \in \mathbb{R}^r$ such that $\|y - x\|_2 = 1$ and $\|y\|_2 = s$. From this, we deduce that

$$[S_x \neq \emptyset \quad \text{if} \quad \|x\|_2 \geq \frac{1}{2}] \quad \text{and} \quad [S'_x \neq \emptyset \quad \text{if} \quad \|x\|_2 \leq 1].$$

Further note that if $0 < \|x\|_2 \leq 1$, then S'_x consists of a single element given by $y = (\frac{\|x\|_2 - 1}{\|x\|_2})x$, while $S'_0 = S^{r-1}$ is the unit sphere.

The Markov chain transitions from $x \in \mathbb{R}^r$ with norm $\|x\|_2 = s$ to the next point according to the following rules.

1. If $0 \leq s < \frac{1}{2}$, move to a point chosen uniformly at random in S'_x .
2. If $\frac{1}{2} \leq s < 1$:
 - (a) with probability $\frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)}$, move to a point chosen uniformly at random in S'_x .
 - (b) with probability $1 - \frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)}$, move to a point chosen uniformly at random in S_x .
3. If $s \geq 1$, then move to a point chosen uniformly at random in S_x .

Symbolically, this defines a transition kernel $q(x, \cdot)$ described explicitly as follows. Let P_x (respectively P'_x) denote the uniform distribution on S_x (respectively S'_x). Then:

$$q(x, \cdot) = \begin{cases} P'_x, & \text{if } 0 \leq s < \frac{1}{2} \\ \frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)} P'_x + (1 - \frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)}) P_x, & \text{if } \frac{1}{2} \leq s < 1 \\ P_x, & \text{if } s \geq 1. \end{cases} \quad (17)$$

This Markov chain is only well-defined if $\chi_{r,\sigma^2}(s) \geq \chi_{r,\sigma^2}(1-s)$ for all $0 \leq s < 1/2$. The following lemma shows that this holds if and only if $\sigma \geq \frac{1}{2\sqrt{r-1}}$.

Lemma 12 *The inequality $\chi_{r,\sigma^2}(s) \geq \chi_{r,\sigma^2}(1-s)$ holds for all $s \in [0, \frac{1}{2}]$ if and only if $\sigma \geq \frac{1}{2\sqrt{r-1}}$.*

Proof Recall that $\chi_{r,\sigma^2}(s) \propto s^{r-1} \exp(-s^2/(2\sigma^2))$, so that the mode is at $\sigma\sqrt{r-1}$ (this can be derived via elementary calculus, since the density is log-concave). In particular, if $\sigma < 1/(2\sqrt{r-1})$, then the mode lies in $(0, 1/2)$, so the desired property cannot hold.

Conversely, suppose $\sigma \geq 1/(2\sqrt{r-1})$. We want to show, for all $s \in [0, 1/2]$,

$$1 \stackrel{!}{\geq} \frac{\chi_{r,\sigma^2}(s)}{\chi_{r,\sigma^2}(1-s)} = \left(\frac{s}{1-s}\right)^{r-1} \exp\left(-\frac{s}{\sigma^2} + \frac{1}{2\sigma^2}\right).$$

Taking logarithms and rearranging, we want

$$\psi(s) := (r-1) \log \frac{s}{1-s} - \frac{s}{\sigma^2} \stackrel{!}{\leq} -\frac{1}{2\sigma^2}.$$

Differentiating, we obtain

$$\psi'(s) \geq (r-1) \left(\frac{1}{s(1-s)} - 4 \right) \geq 0,$$

so ψ is increasing. Hence,

$$\sup_{[0,1/2]} \psi = \psi\left(\frac{1}{2}\right) = -\frac{1}{2\sigma^2}$$

as desired. ■

With this result in hand, we verify that the Markov chain construction has the correct stationary distribution.

Proposition 13 *The Markov transition kernel q on \mathbb{R}^r defines a Markov chain with unit length increments and stationary distribution $\mathcal{N}(0, \sigma^2 I_r)$ provided that $\sigma \geq \frac{1}{2\sqrt{r-1}}$.*

Proof We represent a point in \mathbb{R}^r in spherical coordinates, i.e., as a pair (s, θ) with $s \geq 0$ and $\theta \in S^{r-1}$. Let $\pi := \chi_{r,\sigma^2} \otimes \lambda$, where λ is the uniform distribution on S^{r-1} . Via a change of variables, it suffices to check that for every bounded measurable function $f : \mathbb{R}_+ \times S^{r-1} \rightarrow \mathbb{R}$, it holds that

$$\iint f(s', \theta') \pi(ds, d\theta) q((s, \theta), (ds', d\theta')) = \int f d\pi.$$

We rewrite the LHS of this equation as

$$\begin{aligned} & \iint f(s', \theta') \pi(ds, d\theta) q((s, \theta), (ds', d\theta')) \\ &= \int_0^\infty \left[\iint f(s', \theta') \lambda(d\theta) q((s, \theta), (ds', d\theta')) \right] \chi_{r,\sigma^2}(s) ds \\ &= \left(\int_0^{1/2} + \int_{1/2}^1 + \int_1^\infty \right) \cdots \chi_{r,\sigma^2}(s) ds = \text{I} + \text{II} + \text{III}, \end{aligned}$$

and examine each of the three terms separately.

For the first term,

$$\text{I} = \int_0^{1/2} \left[\iint f(1-s, \theta') \lambda(d\theta) T_{\#} \mathbb{P}'_{(s,\theta)}(d\theta') \right] \chi_{r,\sigma^2}(s) ds,$$

where $T : \mathbb{R}^r \rightarrow S^{r-1}$ is the map $x \mapsto x/\|x\|_2$. Due to the rotational symmetry we have $\int \lambda(d\theta) T_{\#} \mathsf{P}'_{(s,\theta)}(\cdot) = \lambda(\cdot)$, so that evaluating the integral over θ yields:

$$= \int_0^{1/2} \left[\int f(1-s, \theta') \lambda(d\theta') \right] \chi_{r,\sigma^2}(s) ds = \int_{1/2}^1 \left[\int f(s, \theta') \lambda(d\theta') \right] \chi_{r,\sigma^2}(1-s) ds.$$

For the second term, a similar argument yields

$$\begin{aligned} \text{II} &= \int_{1/2}^1 \left[\frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)} \int \int f(1-s, \theta') \lambda(d\theta) T_{\#} \mathsf{P}'_{(s,\theta)}(d\theta') \right. \\ &\quad \left. + \left(1 - \frac{\chi_{r,\sigma^2}(1-s)}{\chi_{r,\sigma^2}(s)}\right) \int \int f(s, \theta') \lambda(d\theta) T_{\#} \mathsf{P}_{(s,\theta)}(d\theta') \right] \chi_{r,\sigma^2}(s) ds \\ &= \int_{1/2}^1 \left[\chi_{r,\sigma^2}(1-s) \int f(1-s, \theta') \lambda(d\theta') + (\chi_{r,\sigma^2}(s) - \chi_{r,\sigma^2}(1-s)) \int f(s, \theta') \lambda(d\theta') \right] ds \\ &= \int_0^{1/2} \left[\int f(s, \theta') \lambda(d\theta') \right] \chi_{r,\sigma^2}(s) ds \\ &\quad + \int_{1/2}^1 \left[\int f(s, \theta') \lambda(d\theta') \right] (\chi_{r,\sigma^2}(s) - \chi_{r,\sigma^2}(1-s)) ds. \end{aligned}$$

Finally, the third term is

$$\text{III} = \int_1^\infty \left[\int \int f(s, \theta') \lambda(d\theta) T_{\#} \mathsf{P}_{(s,\theta)}(d\theta') \right] \chi_{r,\sigma^2}(s) ds = \int_1^\infty \left[\int f(s, \theta') \lambda(d\theta') \right] \chi_{r,\sigma^2}(s) ds.$$

Putting it together, we obtain

$$\text{I} + \text{II} + \text{III} = \int_0^\infty \left[\int f(s, \theta') \lambda(d\theta') \right] \chi_{r,\sigma^2}(s) ds = \int f d\pi$$

which verifies the stationarity of π . ■

For the details of how to sample from the Markov kernel (17) see Appendix C.

3.3. Algorithm and analysis

We develop some notation needed for our algorithm. Let $\sigma_{*,r} = 1/(2\sqrt{r-1})$; recall that this is the smallest standard deviation parameter so that for $\sigma^2 \geq \sigma_{*,r}^2$, the transition probabilities (17) are well-defined. For $\sigma^2 \geq \sigma_{*,r}^2$, we let q_{σ^2} denote the corresponding Markov transition kernel with stationary distribution $\mathcal{N}(0, \sigma^2 I_r)$ as shown in Proposition 13. Also, we let $\mathcal{N}(0, \sigma^2 I_{m \times r})$ denote the law of an $m \times r$ matrix whose entries are i.i.d. $\mathcal{N}(0, \sigma^2)$.

Algorithm 1: Rank- r Gaussian Fixed-Point Walk

Input $v_1, \dots, v_T \in \mathbb{R}^m$ of ℓ_2 -norm at most 1

Initialize $W_0 \sim \mathcal{N}(0, \sigma_{\star,r}^2 I_{m \times r})$
for $t = 1, \dots, T$ **do**

Define $V_t^{(k)} \in \mathbb{R}^{m \times r}$ to be the matrix whose k -th column is v_t and remaining entries are 0.
 $\sigma_t \leftarrow \sigma_{\star,r} / \|v_t\|_2$
 $z_t \leftarrow \frac{1}{\|v_t\|_2^2} W_{t-1}^\top v_t = \frac{1}{\|v_t\|_2^2} (\langle W_{t-1}, V_t^{(k)} \rangle)_{k \in [r]}$
 $z'_t \leftarrow \text{sample from } q_{\sigma_t^2}(z_t, \cdot)$
 $u_t \leftarrow z'_t - z_t$
 $W_t \leftarrow W_{t-1} + v_t u_t^\top = W_{t-1} + \sum_{k=1}^r u_t[k] V_t^{(k)}$

Output $u_1, \dots, u_T \in \mathbb{R}^r$

Proposition 14 *If $\max_{t \in [T]} \|v_t\|_2 \leq 1$, then for all $t \in [T]$, the partial sum $W_t \in \mathbb{R}^{m \times r}$ of Algorithm 1 is distributed as $\mathcal{N}(0, \sigma_{\star,r}^2 I_{m \times r})$.*

Proof We prove this by induction, so suppose that $W_{t-1} \sim \mathcal{N}(0, \sigma_{\star,r}^2 I_{m \times r})$. Let S_t be the subspace $S_t = \text{span}\{V_t^{(k)} : k \in [r]\}$ and observe that

$$W_{t-1} = \text{proj}_{S_t^\perp} W_{t-1} + \text{proj}_{S_t} W_{t-1} = \text{proj}_{S_t^\perp} W_{t-1} + \sum_{k=1}^r z_t[k] V_t^{(k)}.$$

Using standard facts about Gaussian random variables, these two components are independent and $z_t \sim \mathcal{N}(0, \sigma_t^2 I_r)$ (recall the definition of $\sigma_t^2 := \sigma_{\star,r}^2 / \|v_t\|_2^2$). Also, the sampling $z'_t \sim q_{\sigma_t^2}(z_t, \cdot)$ is well-defined because $\sigma_t \geq \sigma_{\star,r}$ by the assumption that $\|v_t\|_2 \leq 1$. Next,

$$W_t = W_{t-1} + \sum_{k=1}^r u_t[k] V_t^{(k)} = \underbrace{\text{proj}_{S_t^\perp} W_{t-1}}_{=: W_t^{(1)}} + \underbrace{\sum_{k=1}^r (z_t + u_t)[k] V_t^{(k)}}_{=: W_t^{(2)}}.$$

Note that $z'_t = z_t + u_t \sim \mathcal{N}(0, \sigma_t^2 I_r)$ and $\|u_t\|_2 = 1$ by Proposition 13. Hence, $W_t^{(2)}$ is distributed as a centered Gaussian on S_t with variance $\sigma_{\star,r}^2$. Moreover, $W_t^{(1)}$ is independent of $W_t^{(2)}$: by the construction of our Markov chain, the distribution of $z'_t = z_t + u_t$ is a function only of z_t , and z_t is independent of $W_t^{(1)}$ as noted above. Thus $W_t \sim \mathcal{N}(0, \sigma_{\star,r}^2 I_{m \times r})$. \blacksquare

The following guarantee for Algorithm 1 is an easy consequence. For a matrix $M \in \mathbb{R}^{m \times r}$, note that the $2 \rightarrow \infty$ norm satisfies $\|M\|_{2 \rightarrow \infty} = \max_{1 \leq i \leq m} \|M_{i,:}\|_2$.

Theorem 15 *If $\max_{t \in [T]} \|v_t\|_2 \leq 1$, then Algorithm 1 outputs $u_1, \dots, u_T \in \mathbb{R}^r$ with the following properties. In expectation,*

$$\mathbb{E} \max_{t \in [T]} \left\| \sum_{s=1}^t v_s u_s^\top \right\|_{2 \rightarrow \infty} \leq \sqrt{\frac{2 \log(mT)}{r-1}} + \sqrt{\frac{r}{r-1}}.$$

Also, with probability at least $1 - \delta$,

$$\max_{t \in [T]} \left\| \sum_{s=1}^t v_s u_s^\top \right\|_{2 \rightarrow \infty} \leq \sqrt{\frac{2 \log(2mT/\delta)}{r-1}} + \sqrt{\frac{r}{r-1}}.$$

Finally, the time complexity of Algorithm 1 is $O(mrT)$.

Proof Note that $\sum_{s=1}^t v_s u_s^\top = W_t - W_0$, where $W_0, W_t \sim \mathcal{N}(0, \sigma_{\star, r}^2 I_{m \times r})$ by Proposition 14, so that $\|\sum_{s=1}^t v_s u_s^\top\|_{2 \rightarrow \infty} \leq \|W_0\|_{2 \rightarrow \infty} + \|W_t\|_{2 \rightarrow \infty}$. Since $\|\cdot\|_{2 \rightarrow \infty}$ is the maximum ℓ_2 norm of a row, then $\max_{t \in [T]} \|W_t\|_{2 \rightarrow \infty}$ can be written as a maximum over mT (not necessarily independent) random variables Y_i , where each Y_i is a norm of a $\mathcal{N}(0, \sigma_{\star, r}^2 I_r)$ random variable. On the other hand, Gaussian concentration of Lipschitz functions (see Theorem 5.5 of Boucheron et al., 2013) implies that each $Y_i \in \text{subG}(\sigma_{\star, r}^2)$, whereas $\mathbb{E} Y_i \leq \sigma_{\star, r} \sqrt{r}$. A similar statement holds for W_0 .

For the expectation bound, a standard sub-Gaussian bound yields

$$\begin{aligned} \mathbb{E} \max_{t \in [T]} \left\| \sum_{s=1}^t v_s u_s^\top \right\|_{2 \rightarrow \infty} &\leq 2 \mathbb{E} \max_{i \in [mT]} Y_i \leq 2 \mathbb{E} \max_{i \in [mT]} (Y_i - \mathbb{E} Y_i) + 2\sigma_{\star, r} \sqrt{r} \\ &\leq 2\sigma_{\star, r} \{ \sqrt{2 \log(mT)} + \sqrt{r} \}, \end{aligned}$$

and the result follows from $\sigma_{\star, r} = \frac{1}{2\sqrt{r-1}}$. The proof of the high-probability bound also uses standard facts about sub-Gaussian random variables and is omitted. ■

As mentioned in Section 1.3, our online vector Komlós result (Theorem 10) is an immediate consequence of Theorem 15. Interestingly, Theorem 15 also implies that scaling $\sigma_{\star, r}^2 \sim \frac{1}{4r}$ of our variance parameter is asymptotically optimal in that no sequence of Markov chains, defined on \mathbb{R}^r for $r \rightarrow \infty$ with unit length increments and Gaussian stationary distribution, can achieve a smaller asymptotic variance. Indeed, a smaller asymptotic variance would lead to an improvement of Theorem 15, which would further yield an improvement of the bound $\text{Vdisc}(A) \leq 1$ of Nikolov (2013). However, the example of the identity matrix shows that this is not possible. We do not know if our variance parameter $\sigma_{\star, r}^2 = \frac{1}{4(r-1)}$ can be improved for any finite r .

We are now ready to prove our main result, the online Banaszczyk bound for Gaussian discrepancy.

Proof [Proof of Theorem 9] Let u_1, \dots, u_T denote the output of Algorithm 1 with rank parameter $r \geq 2$, and let $\Sigma_{i,j}^{(T)} = \langle u_i, u_j \rangle_{\ell_2}$ denote the corresponding Gram matrix. As a consequence of Theorem 15, we have

$$\max_{t \in [T]} \left\| \sum_{s=1}^t v_s u_s^\top \right\|_{2 \rightarrow \infty} \lesssim 1 \vee \sqrt{\frac{\log(mT/\delta)}{r}} \quad (18)$$

with probability at least $1 - \delta$. If $r \geq \log(mT/\delta)$, then combining (18) with (11) from the proof of Proposition 7 implies that

$$\text{Gdisc}(v_1, \dots, v_T; \Sigma^{(T)}) \lesssim \sqrt{\log(2m)}$$

with probability at least $1 - \delta$. If $r \leq \log(mT/\delta)$, then combining (18) with (10) from the proof of Proposition 6 implies that

$$\text{Gdisc}(v_1, \dots, v_T; \Sigma^{(T)}) \lesssim \sqrt{r} \cdot \sqrt{\log(mT/\delta)/r} = \sqrt{\log(mT/\delta)}$$

with probability at least $1 - \delta$. Combining the two cases concludes the proof. ■

Appendix A. EQUIVALENCE OF ONLINE DISCREPANCY FORMULATIONS

In Section 1.3.2 we introduced the online Gaussian discrepancy problem, and described two formulations: one in terms of correlation matrices, and one in terms of unit vectors in ℓ_2 . We now show that the two approaches are equivalent, so that the description of Problem 1 is without loss of generality.

It will be useful to define the *Cholesky decomposition* of a PSD matrix $\Sigma \in \mathbb{R}^{n \times n}$. We define $\text{Ch}(\Sigma)$ to be the unique matrix $L \in \mathbb{R}^{n \times n}$ that is lower triangular, has exactly $r = \text{rank}(\Sigma)$ positive diagonal elements, has $n - r$ columns consisting entirely of zeros, and satisfies

$$\Sigma = LL^\top.$$

We refer the reader to [Gentle \(2012, Chapter 3\)](#) for details regarding existence and uniqueness.

Proposition 16 *The following hold.*

1. *Let $(\Sigma^{(t)})_{t \geq 1}$ be a stream of consistent (in the sense of (13)) correlation matrices. Then one can output unit vectors $(u_t)_{t \geq 1}$ in ℓ_2 online, such that $\Sigma^{(t)}$ is the Gram matrix of u_1, \dots, u_t for each $t \geq 1$. Moreover, only the first t coordinates of u_t need be non-zero.*
2. *Let $(u_t)_{t \geq 1}$ be a stream of unit vectors in ℓ_2 . Then one can output consistent (in the sense of (13)) correlation matrices $(\Sigma^{(t)})_{t \geq 1}$ online, such that $\Sigma^{(t)}$ is the Gram matrix of u_1, \dots, u_t for each $t \geq 1$.*

Proof Note that the second assertion is immediate after defining $\Sigma_{i,j}^{(t)} = \langle u_i, u_j \rangle_{\ell_2}$ for all $i, j \in [t]$ and $t \geq 1$. Now, we turn to the proof of the first assertion. For convenience, we will write ι_n for the inclusion map from \mathbb{R}^k to \mathbb{R}^n where $k \leq n$, with action $\iota_n(x_1, \dots, x_k) = (x_1, \dots, x_k, 0, \dots, 0)$. Let $\tilde{u}_t = \text{Ch}(\Sigma^{(t)})_{t,:} \in \mathbb{R}^t$ be the last row of $\text{Ch}(\Sigma^{(t)})$ for each $t \geq 1$, and let $u_t = \iota_\infty(\tilde{u}_t)$. We claim that the sequence $(u_t)_{t \geq 1}$ satisfies the requirements of the assertion. Clearly only the first t coordinates of u_t can be nonzero. Thus, it remains to show that $\langle u_i, u_j \rangle_{\ell_2} = \Sigma_{i,j}^{(t)}$ for all $i, j \in [t]$, $t \geq 1$.

Define $L^{(t)} = \text{Ch}(\Sigma^{(t)})$ and $r_t = \text{rank}(\Sigma^{(t)})$. It is sufficient to show that for all $t \geq 1$, the rows of $L^{(t)}$ are $\iota_t(\tilde{u}_1), \dots, \iota_t(\tilde{u}_t)$. This is true for $t = 1$ by construction. We proceed by induction and suppose that the claim is true for round $t - 1$. Let L' denote the principal minor of $L^{(t)}$ corresponding to the first $t - 1$ indices. It suffices to show that $L' = L^{(t-1)}$ by the inductive assumption.

Both $L^{(t-1)}$ and L' are lower triangular by construction. Next, by consistency and the definition of the Cholesky decomposition, we have

$$L^{(t-1)} (L^{(t-1)})^\top = \Sigma^{(t-1)} = \Sigma_{[t-1] \times [t-1]}^{(t)} = L' (L')^\top.$$

Therefore, $\text{rank}(L^{(t-1)}) = \text{rank}(L') = r_{t-1}$. Since columns of $L^{(t)}$ with non-positive diagonal elements contain all zeros, it follows that L' has r_{t-1} positive elements on its diagonal. Consequently, L' also has $t - 1 - r_{t-1}$ columns consisting of all zeros. We conclude that $L' = \text{Ch}(\Sigma^{(t-1)}) = L^{(t-1)}$, as desired. ■

Appendix B. FAILURE OF SOME ROUNDING SCHEMES

As a preliminary, we refer to a univariate function $f : \mathbb{R} \rightarrow \mathbb{R}$ as a *sign function* if

$$f(x) = \begin{cases} 1 & \text{if } x > 0 \\ -1 & \text{if } x < 0 \end{cases}$$

and $f(0) \in \{-1, 0, 1\}$. Given a vector $\xi \in \mathbb{R}^n$, we let $\text{sgn}(\xi)$ denote a function given by the entrywise application of sign functions to the coordinates of ξ .

Let $\Sigma \in \mathcal{E}_n$ be the correlation matrix achieving the optimal Gaussian discrepancy. In this section, we consider two ways of rounding Σ to a signing. The Goemans–Williamson rounding ([Goemans and Williamson, 1995](#)) is defined by drawing a random vector $\xi \sim \mathcal{N}(0, \Sigma)$ and setting

$$\sigma_{\text{GW}}(\Sigma) := \text{sgn}(\xi).$$

The PCA rounding is defined by taking a top eigenvector $v_1(\Sigma)$ of Σ and setting

$$\sigma_{\text{PCA}}(\Sigma) := \text{sgn } v_1(\Sigma).$$

Note that when $\Sigma = \sigma\sigma^T$ has rank 1, then $\sigma \in \{\pm 1\}^n$ and both of these rounding schemes recover σ (up to a global sign flip).

In this section, we show that in the setting of both Spencer’s theorem and the Komlós conjecture, these rounding schemes are not very effective. Concretely, for both of these settings, we construct a random matrix that has a planted Gaussian discrepancy solution with objective value zero. However, rounding this solution with either Goemans–Williamson or PCA rounding results in a value of the combinatorial discrepancy that is much larger than the guarantees obtained by direct application of Spencer’s theorem or Banaszczyk’s theorem.

First let us construct our planted Gaussian discrepancy solution. For $n \equiv 2 \pmod{4}$ let U denote the $n \times 2$ matrix whose j -th row is $u_j = (\cos(\frac{2\pi j}{n}), \sin(\frac{2\pi j}{n}))$, and let $\Sigma = UU^T$. Then,

$$\Sigma = \mathbf{c}\mathbf{c}^T + \mathbf{s}\mathbf{s}^T, \tag{19}$$

where $\mathbf{c}_j := \cos(\frac{2\pi j}{n})$ and $\mathbf{s}_j := \sin(\frac{2\pi j}{n})$. Given a vector $v \in \mathbb{R}^n$, define the shift operator $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$ by $s(v) = (v_2, \dots, v_n, v_1)$ and let $w = (\mathbf{1}_{n/2}, -\mathbf{1}_{n/2})$. The lemma below verifies that the Goemans–Williamson and PCA roundings of Σ are close to a vector in $\{w, s(w), \dots, s^{n-1}(w)\}$.

Lemma 17 *For Σ defined as in (19) and $\sigma \in \{\sigma_{\text{GW}}(\Sigma), \sigma_{\text{PCA}}(\Sigma)\}$, the following event holds with probability 1: there exists $k \in \{0, \dots, n-1\}$, $i \neq j \in [n]$, and $a, b \in \{-2, -1, 0, 1, 2\}$ such that*

$$s^k(w) - \sigma = ae_i + be_j.$$

Proof For the Goemans–Williamson rounding, in fact it holds that

$$\mathbb{P}[\sigma_{\text{GW}}(\Sigma) \in \{w, s(w), \dots, s^{n-1}(w)\}] = 1.$$

To see this, observe that $\xi = Ug$ where $g \in \mathbb{R}^2$ is distributed as standard normal. Thus $\sigma_{\text{GW}}(\Sigma)_j = 1$ iff $\langle g, u_j \rangle > 0$. Since the points u_j are equidistributed on the circle S^1 , the claim follows by rotational invariance of the standard normal.

Let us now turn to the PCA-rounding scheme. Since $\mathbf{c} \perp \mathbf{s}$, the top eigenspace (corresponding to eigenvalue $n/2$) of Σ is $\text{span}\{\mathbf{c}, \mathbf{s}\}$. Without loss of generality suppose that our rounding scheme chooses

$$v_1(\Sigma) = \frac{a}{\sqrt{a^2 + b^2}} \mathbf{c} + \frac{b}{\sqrt{a^2 + b^2}} \mathbf{s}$$

as the top eigenvector of Σ for some $(a, b) \neq (0, 0)$. Let θ be such that

$$(\cos \theta, \sin \theta) = \frac{1}{\sqrt{a^2 + b^2}} (a, b).$$

Then we have

$$v_1(\Sigma) = \left(\cos\left(\theta - \frac{2\pi j}{n}\right) \right)_{j \in [n]}.$$

Let $\sigma = \sigma_{\text{PCA}}(\Sigma) = \text{sgn}(v_1(\Sigma))$. We consider two cases. First suppose that $v_1(\Sigma)_j = 0$ for some $j \in [n]$. Then it holds that $\theta \in \frac{\pi}{2} + \frac{2\pi}{n} \mathbb{Z}$. Let $j^* \in [n]$ denote the smallest positive integer such that $\theta + \frac{2\pi j^*}{n} \in \frac{\pi}{2} + \pi \mathbb{Z}$. Recalling that n is even, we have

$$\sigma_j = \begin{cases} \pm 1 & \text{if } j \in (j^*, j^* + n/2) \\ \mp 1 & \text{if } j \in [n] \setminus [j^*, j^* + n/2], \end{cases}$$

and $\sigma_j \in \{-1, 0, 1\}$ if $j \in \{j^*, j^* + n/2\}$. Here we use the properties of the cosine and sign functions. By inspection, the conclusion of the lemma holds in this case.

If on the other hand, $v_1(\Sigma)_j \neq 0$ for all $j \in [n]$, then in fact we have

$$\sigma \in \{w, s(w), \dots, s^{n-1}(w)\},$$

by the fact that n is even and the properties of the cosine and sign functions. The lemma follows. ■

Next, we construct a family of distributions $\mathbf{p}_{m,n}$ on $m \times n$ matrices such that the feasible Gaussian coupling $g \sim \mathcal{N}(0, \Sigma)$ attains zero Gaussian discrepancy. Let $\Sigma^\perp = I_n - \frac{1}{\|\mathbf{c}\|_2^2} \mathbf{c}\mathbf{c}^\top - \frac{1}{\|\mathbf{s}\|_2^2} \mathbf{s}\mathbf{s}^\top$ and write $A \sim \mathbf{p}_{m,n}$ for the random $m \times n$ matrix with rows

$$g_1, \dots, g_m \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma^\perp).$$

The rows of A are isotropic Gaussians on the orthogonal complement of $\text{span}\{\mathbf{c}, \mathbf{s}\}$. Thus in particular $AU = 0$ almost surely so that $\text{Vdisc}_2(A) = 0$ and $\text{Gdisc}(A) = 0$ almost surely.

We prove an intermediate lemma that lower bounds the combinatorial discrepancy incurred by $\sigma \in \{w, s(w), \dots, s^{n-1}(w)\}$.

Lemma 18 *If $A \sim \mathbf{p}_{m,n}$, then there exists a constant $c > 0$ such that for n sufficiently large and for all $\tau \geq 2$,*

$$\mathbb{P} \left[\min_{0 \leq k \leq n-1} \|As^k(w)\|_\infty \leq \tau \right] \leq n \exp \left[-m \exp \left(-\frac{c\tau^2}{n} \right) \right].$$

Proof We record some elementary calculations that will be useful later:

$$\sum_{j=1}^n \cos^2\left(\frac{2\pi j}{n}\right) = \sum_{j=1}^n \sin^2\left(\frac{2\pi j}{n}\right) = \frac{n}{2} \quad (20)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^{n/2} \cos\left(\frac{2\pi j}{n}\right) = 0 \quad (21)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^{n/2} \sin\left(\frac{2\pi j}{n}\right) = \frac{1}{\pi} \quad (22)$$

The identity in the first line is an elementary fact from complex variables, and the limits above are a consequence of Riemann integration.

Observe that for all $1 \leq k \leq n-1$,

$$\left\| \sum_{i=1}^n w_i u_i^\top \right\|_2 = \left\| \sum_{i=1}^n s^k(w)_i u_i^\top \right\|_2$$

by the unitary invariance of the ℓ_2 -norm. Therefore, for all $1 \leq i \leq m$ and $1 \leq k \leq n-1$,

$$\langle g_i, s^k(w) \rangle \sim \mathcal{N}(0, w^\top \Sigma^\perp w).$$

Next,

$$\frac{1}{\|c\|_2^2} \langle c, w \rangle^2 = \frac{1}{n/2} \left\{ 2 \sum_{k=1}^{n/2} \cos\left(\frac{2\pi k}{n}\right) \right\}^2 = o(n), \quad (23)$$

$$\frac{1}{\|s\|_2^2} \langle s, w \rangle^2 = \frac{1}{n/2} \left\{ 2 \sum_{k=1}^{n/2} \sin\left(\frac{2\pi k}{n}\right) \right\}^2 = \frac{8n}{\pi^2} + o(n), \quad (24)$$

by (20), (21), and (22). We have that

$$w^\top \Sigma^\perp w = \underbrace{\left(1 - \frac{8}{\pi^2}\right)}_{=:c_0 > 0} n + o(n).$$

Thus the inner product of a row g_i of A with an output of either Goemans–Williamson or PCA rounding of Σ is a centered Gaussian with variance roughly $c_0 n$.

Next we lower bound $\|A\sigma\|_\infty$ for $\sigma \in \{w, s(w), \dots, s^{n-1}(w)\}$ using the first moment method. Define the random variable

$$S = \sum_{k=0}^{n-1} \mathbb{1}\{\|As^k(w)\|_\infty \leq \tau\}$$

that counts the number of shifts of w attaining discrepancy τ . The Markov inequality implies that

$$\mathbb{P}[\exists k : \|As^k(w)\|_\infty \leq \tau] = \mathbb{P}[S \geq 1] \leq \mathbb{E} S. \quad (25)$$

By our work above, all of the vectors $As^k(w)$ have the same distribution and have independent coordinates. Using this fact and standard Gaussian estimates, for $\tau \geq 2$,

$$\begin{aligned}\mathbb{E} S &= n \mathbb{P}[\|Aw\|_\infty \leq \tau] = n \mathbb{P}_{g \sim \mathcal{N}(0, w^\top \Sigma^\perp w)}[|g| \leq \tau]^m \\ &\leq n \left\{ 1 - \exp\left(-\frac{\tau^2}{2 \langle w, \Sigma^\perp w \rangle}\right) \right\}^m \\ &\leq n \exp\left[-m \exp\left(-\frac{\tau^2}{2(c_0 + o(1))n}\right)\right].\end{aligned}\tag{26}$$

The lemma statement follows from (25) and the previous display. \blacksquare

Lemma 18 allows us to demonstrate the poor performance of the rounding schemes considered in this section in two well-studied settings.

Proposition 19 *Let $A \sim \mathfrak{p}_{m,n}$ and define $A' = \frac{1}{c\sqrt{\log n}} A$ for a sufficiently large absolute constant $c > 0$. Let Σ be defined as in (19), and let $\sigma \in \{\sigma_{\text{GW}}, \sigma_{\text{PCA}}\}$. Then*

$$\text{Gdisc}(A') = 0 = \mathbb{E}\|A'g\|_\infty \quad \text{if } g \sim \mathcal{N}(0, \Sigma),$$

and the following hold with high probability as $n \rightarrow \infty$:

- (Spencer setting) If $m = \frac{n}{\log n}$, then $\max_{i \in [m], j \in [n]} |A'_{i,j}| \leq 1$, and

$$\text{disc}(A') \lesssim \sqrt{\frac{n}{\log n}} \ll \sqrt{n} \lesssim \|A'\sigma\|_\infty.$$

- (Komlós setting) If $m = 10 \log n$, then $\max_{j \in [n]} \|A'_{:,j}\|_2 \leq 1$, and

$$\text{disc}(A') \lesssim \sqrt{\log \log n} \ll \sqrt{\frac{n}{\log n}} \lesssim \|A'\sigma\|_\infty.$$

Proof Recall that $\mathbb{E}\|A'g\|_\infty = 0$ holds for $g \sim \mathcal{N}(0, \Sigma)$ by construction.

Next we verify the high probability bounds on the size of the entries of A' . The bound on the size of the largest entry of A' in the Spencer setting follows from the union bound and tail bounds for the standard Gaussian distribution. Now, for the Komlós setting, since $\Sigma^\perp \preceq I_n$, it follows that $\text{var}(A_{i,j}) \leq 1$ for all i, j . Let $\tau_j \leq 1$ denote the common variance of the entries in column j . We have that $\|A_{:,j}\|_2 - \sqrt{\tau_j m}$ is sub-Gaussian with variance proxy 1 (see e.g. [Vershynin, 2018](#), Chapter 3). Applying the sub-Gaussian tail bound and the union bound over $1 \leq j \leq n$ implies that $\max_{j \in [n]} \|A_{:,j}\|_2 \lesssim \sqrt{\log n}$, as desired.

Given the previous bounds on the entries of A' , the upper bounds on $\text{disc}(A')$ follow immediately from Spencer's theorem and Banaszczyk's bound, respectively. To finish, we verify the lower bounds on $\|A'\sigma\|_\infty$ when $\sigma \in \{\sigma_{\text{GW}}, \sigma_{\text{PCA}}\}$. By Lemma 17, it holds with probability 1 that for $\sigma \in \{\sigma_{\text{GW}}, \sigma_{\text{PCA}}\}$,

$$s^k(w) - \sigma = ae_i + be_j$$

for some $k \in \{0, \dots, n-1\}$, $i \neq j \in [n]$, and $|a|, |b| \leq 2$. By the union bound over this event and the event that $\max_{i,j} |A'_{i,j}| \leq 1$, it holds with high probability that

$$\|A'(\sigma - s^k(w))\|_\infty \leq 4,\tag{27}$$

for some $k \in \{0, \dots, n-1\}$, by the triangle inequality. By the union bound again over the event (27) and the event of Lemma 18, it holds that

$$\|A'\sigma\|_\infty \geq \min_{0 \leq k \leq n-1} \|A's^k(w)\|_\infty - 4 \gtrsim \begin{cases} \sqrt{n} & \text{if } m = \frac{n}{\log n} \\ \sqrt{\frac{n}{\log n}} & \text{if } m = 10 \log n \end{cases}.$$

Precisely, for $m = \frac{n}{\log n}$, we set $\tau = c'\sqrt{n \log n}$ in Lemma 18, and for $m = 10 \log n$, we set $\tau = c''\sqrt{n}$, for sufficiently small absolute constants $c', c'' > 0$. ■

Furthermore, when $m = \frac{n}{\log n}$, a signing chosen uniformly at random already achieves discrepancy $O(\sqrt{n})$ for the matrix A' with high probability. This indicates that in the Spencer setting our rounding schemes σ_{GW} and σ_{PCA} fail to ‘beat the union bound’. We justify this below for completeness.

Lemma 20 *With high probability over the random matrix $A \sim \mathfrak{p}_{m,n}$, a uniformly random sign $\sigma \sim \text{unif}(\{\pm 1\}^n)$ satisfies*

$$\mathbb{E}[\|A'\sigma\|_\infty \mid A'] \lesssim \sqrt{n},$$

where $A' = \frac{1}{c\sqrt{\log n}}A$ as before.

Proof By standard concentration results (see, e.g., Boucheron et al., 2013, Theorem 5.5) and using the fact that $\text{var}(A_{i,j}) \leq 1$ for all i, j , it holds for fixed i that

$$\mathbb{P}[\|A_{i,:}\|_2 - \sqrt{n} \geq t] \leq \exp(-t^2/2).$$

Thus the union bound implies that the event $E = \{\max_{1 \leq i \leq m} \|A_{i,:}\|_2 \leq 2\sqrt{n}\}$ holds with probability at least $1 - me^{-n/2}$. Conditionally on the matrix A , the random variable $\sum_{j=1}^n A_{i,j}\sigma_j$ is sub-Gaussian with parameter $\|A_{i,:}\|_2^2$, so the maximal inequality for sub-Gaussian random variables (see, e.g., Boucheron et al., 2013, Theorem 2.5) yields

$$\mathbb{E}[\|A\sigma\|_\infty \mid A] \leq \left(\max_{1 \leq i \leq m} \|A_{i,:}\|_2 \right) \sqrt{2 \log(2m)} \leq 2 \sqrt{2n \log \frac{2n}{\log n}} = O(\sqrt{n \log n})$$

with probability at least $1 - me^{-n/2}$, as desired. ■

Appendix C. SAMPLING FROM THE MARKOV CHAIN

To implement the Gaussian fixed point walk in dimension $r > 1$, we need to be able to sample from the Markov kernel (17). Clearly, to this end it is sufficient to solve the following problem.

Problem 2 *Given a point $x \in \mathbb{R}^r$ and a target radius $s \geq 0$, output a uniformly random point from*

$$S := \{y \in \mathbb{R}^r : \|x - y\|_2 = 1, \|y\|_2 = s\}, \quad (28)$$

provided $S \neq \emptyset$.

Let $\delta = y - x$. Equivalently, we need to sample from the set of $\delta \in \mathbb{R}^r$ with $\|\delta\|_2 = 1$ and $\|x + \delta\|_2 = s$. Expanding and rearranging the latter, we obtain

$$\left\langle \frac{x}{\|x\|_2}, \delta \right\rangle = \frac{s^2 - \|x\|_2^2 - 1}{2\|x\|_2}.$$

The set of such δ can be parametrized as $\delta = \lambda x/\|x\|_2 + \sqrt{1 - \lambda^2} w$ where w is a unit vector orthogonal to x , and $\lambda = (s^2 - \|x\|_2^2 - 1)/(2\|x\|_2)$. Uniformly sampling such a w is easy: given a standard Gaussian $g \sim \mathcal{N}(0, I_r)$, project it onto x^\perp and normalize:

$$w = \frac{(I_r - xx^\top/\|x\|_2) g}{\|(I_r - xx^\top/\|x\|_2) g\|_2}. \quad (29)$$

Putting it all together, we obtain the following algorithm.

Algorithm 2: Sampling from the Markov kernel (17)

Input: $x \in \mathbb{R}^r$, $s \geq 0$.
Set $\lambda = (s^2 - \|x\|_2^2 - 1)/(2\|x\|_2)$.
if $\lambda \notin [0, 1]$ **then**
 └ Output: S is empty.
else
 └ Draw w according to (29).
 └ Set $\delta = \lambda x/\|x\|_2 + \sqrt{1 - \lambda^2} w$.
 └ Output: $x + \delta$.

ACKNOWLEDGMENTS

We thank Kwangjun Ahn and Chen Lu for helpful conversations. SC was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. PR is supported by NSF grants IIS-1838071, DMS-2022448, and CCF-2106377.

References

- D. Alonso-Gutiérrez and J. Bastero. *Approaching the Kannan–Lovász–Simonovits and variance conjectures*, volume 2131 of *Lecture Notes in Mathematics*. Springer, Cham, 2015.
- J. M. Altschuler and E. Boix-Adserà. Hardness results for multimarginal optimal transport problems. *Discrete Optim.*, 42:Paper No. 100669, 21, 2021.
- R. Alweiss, Y. P. Liu, and M. Sawhney. *Discrepancy minimization via a self-balancing walk*, page 14–20. Association for Computing Machinery, New York, NY, USA, 2021.
- W. Banaszczyk. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Structures and Algorithms*, 12:351–360, 1998.
- N. Bansal. Constructive algorithms for discrepancy minimization. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS ’10, pages 3–10, Washington, DC, USA, 2010. IEEE Computer Society.

N. Bansal and J. Spencer. On-line balancing of random inputs. *Random Structures & Algorithms*, 57(4):879–891, 2020.

N. Bansal, D. Dadush, and S. Garg. An algorithm for Komlós conjecture matching Banaszczyk’s bound. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 788–799, 2016.

N. Bansal, D. Dadush, S. Garg, and S. Lovett. The Gram-Schmidt walk: a cure for the Banaszczyk blues. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 587–597, 2018.

N. Bansal, H. Jiang, S. Singla, and M. Sinha. Online vector balancing and geometric discrepancy. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1139–1152, 2020.

N. Bansal, H. Jiang, R. Meka, S. Singla, and M. Sinha. Online discrepancy minimization for stochastic arrivals. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2842–2861. SIAM, 2021.

J. Beck and T. Fiala. “Integer-making” theorems. *Discrete Appl. Math.*, 3(1):1–8, 1981.

S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.

M. Charikar, A. Newman, and A. Nikolov. Tight hardness results for minimizing discrepancy. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1607–1614. SIAM, 2011.

B. Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, Cambridge, 2000.

K. Costello. Balancing Gaussian vectors. *Israel Journal of Mathematics*, 172(1):145–156, 2009.

D. Dadush, S. Garg, S. Lovett, and A. Nikolov. Towards a constructive version of Banaszczyk’s vector balancing theorem. *Theory of Computing*, 15(1):1–58, 2019.

R. Eldan and M. Singh. Efficient algorithms for discrepancy minimization in convex sets. *Random Structures & Algorithms*, 53(2):289–307, 2018.

J. E. Gentle. *Numerical linear algebra for applications in statistics*. Springer Science & Business Media, 2012.

A. A. Giannopoulos. On some vector balancing problems. *Studia Math.*, 122(3):225–234, 1997.

E. D. Gluskin. Extremal properties of orthogonal parallelepipeds and their applications to the geometry of Banach spaces. *Mathematics of the USSR-Sbornik*, 64(1):85, 1989.

M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

C. Harshaw, F. Sävje, D. Spielman, and P. Zhang. Balancing covariates in randomized experiments with the Gram–Schmidt walk design. *arXiv e-prints*, art. arXiv:1911.03071, 2021.

N. J. A. Harvey, R. Schwartz, and M. Singh. Discrepancy without partial colorings. In *Approximation, randomization, and combinatorial optimization*, volume 28 of *LIPICS. Leibniz Int. Proc. Inform.*, pages 258–273. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2014.

C. Jones and M. McPartlon. Spherical discrepancy minimization and algorithmic lower bounds for covering the sphere. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 874–891. SIAM, 2020.

N. Karmarkar, R. Karp, G. Lueker, and A. Odlyzko. Probabilistic analysis of optimum partitioning. *Journal of Applied Probability*, 23:626–645, September 1986.

A. M. Krieger, D. Azriel, and A. Kapelner. Nearly random designs with greatly improved balance. *Biometrika*, 106(3):695–701, 2019.

A. Levy, H. Ramadas, and T. Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, pages 380–391, 2017.

Y. P. Liu, A. Sah, and M. Sawhney. A Gaussian fixed point random walk. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPICS)*, pages 101:1–101:10, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS ’12*, pages 61–67, Washington, DC, USA, 2012. IEEE Computer Society.

J. Matoušek. *Geometric discrepancy: an illustrated guide*. Springer, New York, 1999.

Y. Nesterov. *Lectures on convex optimization*, volume 137 of *Springer Optimization and Its Applications*. Springer, Cham, 2018.

A. Nikolov. The Komlós conjecture holds for vector colorings. *arXiv e-prints*, art. arXiv:1301.4039, 2013.

B. Pass. Multi-marginal optimal transport: theory and applications. *ESAIM Math. Model. Numer. Anal.*, 49(6):1771–1790, 2015.

T. Rothvoss. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.*, 46(1):224–234, 2017.

J. Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289:679–706, 1985.

P. Turner, R. Meka, and P. Rigollet. Balancing Gaussian vectors in high dimension. In *Conference on Learning Theory*, pages 3455–3486. PMLR, 2020.

R. Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. An introduction with applications in data science, With a foreword by Sara van de Geer.