Vector Balancing and Integer Programming

Victor Reis

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2023

Reading Committee:

Thomas Rothvoss, Chair

Anup Rao

Kevin Jamieson

Program Authorized to Offer Degree:

Computer Science & Engineering

University of Washington

Abstract

Vector Balancing and Integer Programming

Victor Reis

Chair of the Supervisory Committee:
Associate Professor Thomas Rothvoss
Mathematics and Computer Science & Engineering

A set of $2n$ pennies can be easily *balanced* into two groups of equal weight, yet doing so by tossing each coin would incur a *standard deviation* of $\Theta(\sqrt{n})$ as an imbalance in the weight of the two groups. More realistically, when running randomized controlled trials for vaccines and deciding who will receive a vaccine and who will receive a placebo, it is desirable to design robust experiments which are also balanced [71]. This is because significant imbalances, say in the average age of each group or in any other relevant attribute, reduce the value of the experiment. Balancing has found other applications in fair resource allocation [105], differential privacy [126], bin packing [74], and scheduling [26].

We start by motivating the study of vector balancing with five fascinating open problems (Chapter 0). We show improved vector balancing bounds for several classes of convex bodies, such as ℓ_p balls, zonotopes and Schatten balls (Chapters 2, 3 and 7). We explore connections to sparsification of convex combinations (Chapter 4) and graphs (Chapter 6). We investigate the tightness of approximations to a robust notion of balancing (Chapter 5), prefix balancing problems (Chapter 8) and online settings (Chapter 9).

Finally, we give a faster algorithm for integer programming (Chapter 10), a fundamental problem in discrete optimization, by providing a constructive answer to a question of Kannan and Lovász [83] and Dadush [58] on the subspace flatness of convex bodies.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

It has been an honor to have Thomas as my advisor. Even before coming to UW, I was struck by the clarity and elegance of his lecture notes, and it was a wonderful surprise to see that his informal research notes are held to a similar standard of precision. I am often amazed by his perseverance towards solving meaningful problems and by the seriousness with which he listens: for example, there was a time during the first month of my Ph.D. when I had briefly mentioned a textbook, and two years later not only had he read it, but also decided to teach it as a graduate course.

I would also like to thank the other members of my Ph.D. supervisory committee, including Anup Rao, Kevin Jamieson, and Zaid Harchaoui. I took Anup's classes in coding theory and communication complexity, and indeed his concise communication style, as well as the feedback he gave me on my presentations, have had a lasting impact on me. Kevin and Zaid devoted a lot of time and effort to supervising my progress since my General Examination in 2022.

I am beholden to Daniel Dadush, not just for the excellent presentation of the integer programming problem in the last chapter of his thesis (which motivated the last chapter of mine), but also for the incredible mentorship he provided me over the course of my Ph.D.. Since the first time he took me as a visiting student for a February 2020 lattices workshop in Berkeley, Daniel has been an inspiring research collaborator, joining evening Zoom calls in spite of the nine hour time difference, and more recently, inviting me to visit him in Amsterdam on May 2023.

I remain deeply grateful to people I met before my Ph.D. who influenced me to pursue further education, including Rodrigo Angelo, Luciano de Castro, Alexandre de Faveri, Daniel Fleischman, Chloe Hsu, Dexter Kozen, Robert Kleinberg, Carlos Shine, Eduardo Tengan and David Williamson.

I would like to thank all my research collaborators and academic friends I met during my Ph.D., including Sander Borst, Rainie Bozzai, Sami Davies, Sally Dong, Siddarth Iyer, Arun Jambulapati, Haotian Jiang, Ce Jin, Nathan Klein, James Lee, Kuikui Liu, Yang Liu, Sasho Nikolov, Swati Padmanabhan, Lars Rohwedder, Zhao Song, Ruoqi Shen, Kevin Tian, Amir Yehudayoff and Xinzhi Zhang. I would also like to thank Sivakanth Gopi, Janardhan Kulkarni and Sergey Yekhanin for hosting me at the Algorithms group at Microsoft Research in Redmond on the summer of 2023.

Special thanks to the people who made lockdowns bearable: to the morning crew who beat me up at Framework Brazilian Jiu-Jitsu; to the UW players who I beat up in Teamfight Tactics; to Ewin Tang, who was my housemate for three years and my best friend in Seattle; and to Doris Yang, who was my partner for four years.

I thank the UW CSE department, in particular Magda Bałazińska and Stuart Reges, for their support towards the annual programming contest I ran for four years during my Ph.D., which turned out to be a wonderful creative outlet at times when research was hard. I also thank the hundreds of students who enthusiastically formed teams to compete. In particular, I send my love to Alicia Stepin, for flooding my life with joy and hope for a lovely future together.

Most of all, I thank my parents Ana Patrícia and Marfran for their unrelenting love and support through all my academic life.

Chapter 0

FIVE OPEN PROBLEMS IN VECTOR BALANCING

0.1 Introduction

The *vector balancing constant* for two symmetric convex bodies $P, Q \subseteq \mathbb{R}^d$, first considered by Dvoretzky in 1963 [60], is defined as

$$\mathsf{vb}(P,Q) := \sup\Big\{ \min_{\boldsymbol{x}\in\{-1,1\}^n} \Big\| \sum_{i=1}^{n} x_i \boldsymbol{v}_i \Big\|_Q \mid n \in \mathbb{N}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P \Big\}.$$

The signs $\boldsymbol{x} \in \{-1,1\}^n$ are often referred to as a (*full*) *coloring* of the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$. We also define

$$\mathsf{vb}_n(P,Q) := \sup\Big\{ \min_{\boldsymbol{x}\in\{-1,1\}^n} \Big\| \sum_{i=1}^{n} x_i \boldsymbol{v}_i \Big\|_Q \mid \boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P \Big\}$$

as the vector balancing constant restricted to n vectors[1], so $\mathsf{vb}(P,Q) = \sup_{n\in\mathbb{N}} \mathsf{vb}_n(P,Q)$. We will see in Chapter 1 that this makes sense as $\mathsf{vb}(P,Q) \leq 2 \cdot \mathsf{vb}_d(P,Q)$. The simplest example of a vector balancing constant is $\mathsf{vb}([-1,1],[-1,1]) \leq 1$ which can be verified by a simple greedy approach, picking signs one by one so as to make the current sum closer to zero.

While the study of vector balancing problems does have several applications, a fundamental motivation in studying them is the simplicity of the statement of the hardest problems in the field, which remained elusive for several decades and spurred the development of several technical tools in convex geometry [13, 146], probability [16, 103], linear algebra [28, 19], Fourier analysis [75, 65] and polynomials [107].

[1] We should note that this is the same as asking for *at most* n vectors as $\mathbf{0} \in P$.

0.2 Five Open Problems

- **The Komlós conjecture.** The simplest case of the vector balancing problem in high dimensions is without a doubt the Euclidean setting, whereby taking random signs shows $\mathsf{vb}_d(B_2^d, B_2^d) \le \sqrt{d}$; a lower bound of $\sqrt{d}$ also follows from any orthonormal basis. In a seminal work, Spencer showed in 1985 that a similar bound holds for the cube, namely[2] $\mathsf{vb}_d(B_\infty^d, B_\infty^d) \le 6\sqrt{d}$ [157], and Walsh-Hadamard matrices show that this bound is tight up to constants. The problem of Komlós, communicated by Spencer in the same work, asks for a common generalization of these two results:

  Conjecture 1. *Does there exist a universal constant $C > 0$ so that $\mathsf{vb}(B_2^d, B_\infty^d) \le C$?*

  A simple construction due to Kunisky shows that $C \ge 1 + \sqrt{2}$ [91] and the best known bound for n vectors is $\mathsf{vb}_n(B_2^d, B_\infty^d) \lesssim \sqrt{\log \min(n, d)}$ [13, 30, 19]. In Chapter 2, we explore related vector balancing constants of the form $\mathsf{vb}(B_p^d, B_q^d)$. In Chapter 9, we provide an online algorithm that achieves the best known bound.

- **The Steinitz problem and signed series.** Perhaps the oldest problem in the vector balancing literature is the Steinitz problem. We define the Steinitz constant $\mathsf{St}_n(P, Q)$ as

  $$\mathsf{St}_n(P,Q) := \sup\Big\{ \min_{\pi \in S_n} \max_{k \in [n]} \Big\| \sum_{i=1}^k \boldsymbol{v}_{\pi(i)} \Big\|_Q \mid \boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P, \sum_{i=1}^n \boldsymbol{v}_i = 0 \Big\}.$$

  This is in fact a stronger notion of vector balancing:

  Lemma 1. [*14*] *For any $n \in \mathbb{N}$, $\mathsf{vb}_n(P, Q) \le 2 \cdot \mathsf{St}_{2n+1}(P, Q)$.*

  Proof. Given $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P$, construct an instance of the Steinitz problem by adding $M := 2\lceil n/2 \rceil$ copies of the vector $\boldsymbol{u} := -\sum_{i=1}^n \boldsymbol{v}_i / M \in P$. The prefix sum which includes exactly $M/2$ copies of $\boldsymbol{u}$ corresponds to (half) signs $\boldsymbol{x} \in \{-1/2, 1/2\}^n$. $\square$

[2] A conjecture of Spielman is that a bound of $2\sqrt{d}$ should also hold. In the lower bound side, an 8×8 Walsh-Hadamard matrix with the first column replaced by zeros shows that the factor of 6 cannot be replaced by a number smaller than $5/\sqrt{8} > 1.76$.

Steinitz showed in 1913 [159] that $\mathsf{St}(K,K) \le 2d$ for any symmetric convex K, and this was improved to d by Grinberg and Sevastyanov [69]. Yet even for the simplest of high-dimensional convex bodies, the value of the Steinitz constant is not known:

Conjecture 2. *Let $K \in \{B_2^d, B_\infty^d\}$. Is it true that $\mathsf{St}(K,K) \lesssim \sqrt{d}$?*

A lower bound of $\Omega(\sqrt{d})$ may be derived from Lemma 1. The best known bounds depending on the number of vectors are due to Banaszczyk [14] who showed that $\mathsf{St}_n(B_2^d, B_2^d) \lesssim \sqrt{d} + \sqrt{\log n}$ and $\mathsf{St}_n(B_\infty^d, B_\infty^d) \lesssim \sqrt{d \log n}$. Another interesting connection to vector balancing was shown by Chobanyan [46], namely the relation $\mathsf{St}_n(P,Q) \le \mathsf{ss}_n(P,Q)$ for any $n \in \mathbb{N}$, where ss is the *signed series* constant

$$\mathsf{ss}_n(P,Q) := \sup\Big\{ \min_{\boldsymbol{x} \in \{-1,1\}^n} \max_{k \in [n]} \Big\| \sum_{i=1}^k x_i \boldsymbol{v}_i \Big\|_Q \mid \boldsymbol{v}_1, \dots, \boldsymbol{v}_n \in P \Big\}.$$

Bárány and Grinberg showed a bound of $\mathsf{ss}_n(K,K) \le 2d-1$ for any symmetric convex K [41, 28] and a bound of $O(\sqrt{d})$ is also conjectured as in Conjecture 2. In Chapter 2, we explore bounds on $\mathsf{ss}_n(B_p^d, B_q^d)$.

Another interesting quantity that allows for better bounds independent of the number of vectors is the *signed rearrangements constant*

$$\mathsf{sr}_n(P,Q) := \sup\Big\{ \min_{\pi \in S_n} \min_{\boldsymbol{x} \in \{-1,1\}^n} \max_{k \in [n]} \Big\| \sum_{i=1}^k x_i \boldsymbol{v}_{\pi(i)} \Big\|_Q \mid \boldsymbol{v}_1, \dots, \boldsymbol{v}_n \in P \Big\}.$$

By taking $k = n$ in the above definition we see $\mathsf{vb}_n(P,Q) \le \mathsf{sr}_n(P,Q)$, and $\pi = \mathrm{id}$ reveals $\mathsf{sr}_n(P,Q) \le \mathsf{ss}_n(P,Q)$. In fact, the signed rearrangements constant may also be related to the Steinitz constant:

Lemma 2. [14] *For any $n \in \mathbb{N}$, $\mathsf{sr}_n(P,Q) \le \mathsf{vb}_n(P,Q) + \mathsf{St}_{2n}(P,Q)$.*

Proof. Given $\boldsymbol{v}_1, \dots, \boldsymbol{v}_n$, let $\boldsymbol{x} \in \{-1,1\}^n$ be signs so that $\sum_{i=1}^n x_i \boldsymbol{v}_i \in \mathsf{vb}_n(P,Q) \cdot Q$. Let $\boldsymbol{u} := -\sum_{i=1}^n x_i \boldsymbol{v}_i / n \in P$ and consider the permutation π' of the set $\{x_1 \boldsymbol{v}_1, \dots, x_n \boldsymbol{v}_n, \boldsymbol{u}, \dots, \boldsymbol{u}\}$ (n copies of $\boldsymbol{u}$) with prefix sums in $\mathsf{St}_{2n}(P,Q) \cdot Q$. Then π' induces a permutation $\pi \in S_n$ so that $\sum_{i=1}^k x_{\pi(i)} \boldsymbol{v}_{\pi(i)} \in (\mathsf{vb}_n(P,Q) + \mathsf{St}_{2n}(P,Q)) \cdot Q$. $\square$

Banaszczyk showed $\mathsf{sr}_n(P,Q) \le 2 \cdot \mathsf{ss}_{\lfloor 8d \log d \rfloor}(P,Q)$ for any $n \in \mathbb{N}$ and any $d \ge 2$, and in particular that $\mathsf{sr}_n(B_2^d, B_2^d) \lesssim \sqrt{d}$ for all $n \in \mathbb{N}$. On the other hand, it remains an open problem to determine $\mathsf{sr}(B_\infty^d, B_\infty^d) := \sup_{n\in\mathbb{N}} \mathsf{sr}_n(B_\infty^d, B_\infty^d)$:

Conjecture 3. *Is it true that* $\mathsf{sr}(B_\infty^d, B_\infty^d) \lesssim \sqrt{d}$*?*

In Chapter 2 we show $\mathsf{sr}(B_\infty^d, B_\infty^d) \lesssim \sqrt{d}\log\log d$, improving upon the $O(\sqrt{d \log d})$ bound of Banaszczyk [14].

- **The vector balancing constant for zonotopes.** As mentioned previously, a seminal result of Spencer [157] is that $\mathsf{vb}(B_\infty^d, B_\infty^d) \lesssim \sqrt{d}$. In Chapter 2 we will see that indeed $\mathsf{vb}(B_p^d, B_p^d) \lesssim \sqrt{d}$ for $2 \le p \le \infty$. A fundamental question raised by Schechtman [152] is whether the same holds for any *zonotope*, the linear image of a cube:

  Conjecture 4. *Let* $K = AB_\infty^m \subseteq \mathbb{R}^d$*, where* $A \in \mathbb{R}^{d\times m}$*. Is it true that* $\mathsf{vb}(K,K) \lesssim \sqrt{d}$*?*

  We refer to Chapter 3, where in particular we show $\mathsf{vb}(K,K) \lesssim \sqrt{d}\log\log\log d$.

- **The matrix Spencer conjecture.** Another extension of Spencer's theorem was asked by Zouzias [171] (see also [113]):

  Conjecture 5. *Given matrices* $A_1, \ldots, A_d \in \mathbb{R}^{d\times d}$ *with bounded operator norm* $\|A_i\|_{\mathrm{op}} \le 1$*, do there exist signs* $x \in \{-1,1\}^d$ *so that* $\|\sum_{i=1}^d x_i A_i\|_{\mathrm{op}} \lesssim \sqrt{d}$*?*

  Here $\|A\|_{\mathrm{op}}$ denotes the largest singular value of A. Note that when the matrices are diagonal, this corresponds to the ℓ_∞ norm of the vector formed by the entries of the diagonal. We refer to Chapter 7 for further discussion, where in particular we show a bound of $\sqrt{n\log(2d^2/n)}$ for n matrices.

Chapter 1

VECTOR BALANCING TOOLBOX

This chapter is based on useful tools from joint papers with Thomas Rothvoss [140], Rainie Bozzai and Thomas Rothvoss [38], Daniel Dadush and Haotian Jiang [52] and Arun Jambulapati and Kevin Tian [77].

On a first read, most of this chapter should be skipped, and referred to as tools prove necessary later on.

1.1 Introduction

We start by justifying the bound from the previous chapter on the number of vectors needed for the vector balancing constant:

Theorem 3. *For any symmetric convex bodies $P, Q \subseteq \mathbb{R}^d$, $\mathsf{vb}(P,Q) \leq 2 \cdot \mathsf{vb}_d(P,Q)$.*

Proof. Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P$ with $n > d$ and consider a basic solution $\boldsymbol{x}^*$ to the linear program $\sum_{i=1}^n x_i \boldsymbol{v}_i = \mathbf{0}$, $-1 \leq x_i \leq 1$ for $i \in [n]$. Let $I := \{i \in [n] : -1 < x_i^* < 1\}$ denote the fractional indices and note that $|I| \leq d$. A theorem of Lovász, Spencer and Vesztergombi [102] allows one to round fractional solutions while incurring an error of at most $2 \cdot \mathsf{vb}_d(P,Q)$. Namely, there exists an $\boldsymbol{x} \in \{-1,1\}^I$ so that

$$\Big\| \sum_{i \in I} (x_i - x_i^*) \boldsymbol{v}_i \Big\|_Q \leq 2 \cdot \mathsf{vb}_d(P,Q).$$

Then we may extend $\boldsymbol{x}$ into a full coloring in $\{-1,1\}^n$ by setting $x_i := x_i^* \in \{-1,1\}$ for $i \notin I$, so that $\| \sum_{i=1}^n x_i \boldsymbol{v}_i \|_Q = \| \sum_{i \in I} (x_i - x_i^*) \boldsymbol{v}_i \|_Q \leq 2 \cdot \mathsf{vb}_d(P,Q)$. $\square$

We have therefore reduced the vector balancing problem to d many vectors. The main technique we will use throughout this dissertation to achieve upper bounds for this case is that of *partial coloring*: it turns out to be easier to find a vector $\boldsymbol{x} \in [-1,1]^n$ with $\Omega(n)$ coordinates in $\{-1,1\}$ and iterate to get a full coloring. When $\boldsymbol{x} \notin \{-1,0,1\}^n$, such a partial coloring is called fractional.

A result of Giannopoulos [67] shows that for a *small enough* constant $\alpha > 0$, a symmetric convex body K with Gaussian measure at least $e^{-\alpha n}$ contains a partial coloring $\boldsymbol{x} \in \{-1,0,1\}^n$ with a linear number of entries in ± 1. In this chapter, we will show an algorithmic version of this result. First, we define Gaussian measure and mention some key properties.

Gaussian Measure. We use $\gamma_n(\cdot)$ to denote the standard Gaussian measure on $\mathbb{R}^n$. Gaussian measure is log-concave, i.e. $\gamma_n(\lambda A+(1-\lambda)B) \geq \gamma_n(A)^\lambda \gamma_n(B)^{1-\lambda}$ for any compact subsets $A, B \subseteq \mathbb{R}^n$. In particular, by taking $A = -x + K$ and $B = x + K$ for any $x \in \mathbb{R}^n$ and symmetric convex body K, and $\lambda = 1/2$, we have the following lemma.

Lemma 4 (Translation Decreases Gaussian Measure)**.** *Given any symmetric convex body $K \subseteq \mathbb{R}^n$ and $x \in \mathbb{R}^n$, we have $\gamma_n(K) \geq \gamma_n(x + K)$.*

A well-known correlation inequality is the following:

Lemma 5 (Šidak [155] and Kathri [86])**.** *For any symmetric convex set $K \subseteq \mathbb{R}^n$ and strip $S = \{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{a}, \boldsymbol{x}\rangle| \leq 1\}$, one has $\gamma_n(K \cap S) \geq \gamma_n(K) \cdot \gamma_n(S)$.*

It is worth noting that a recent result of Royen [147] extends this to any two arbitrary symmetric sets. We refer to the exposition of Latała and Matlak [93]:

Theorem 6 (Gaussian Correlation Inequality)**.** *Given any symmetric convex sets $K, T \subseteq \mathbb{R}^n$, we have $\gamma_n(K \cap T) \geq \gamma_n(K) \cdot \gamma_n(T)$.*

The main tool we use to find partial colorings is a constructive version of Giannopoulos' result which shows that for fractional colorings *any* constant $\alpha > 0$ suffices. Our argument even works for intersections with a large enough subspace.

Theorem 7. *For all $\alpha, \beta, \gamma > 0$, there is a constant $C := C(\alpha, \beta, \gamma) > 0$ so that the following holds: There is a randomized polynomial time algorithm which for a symmetric convex set $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq e^{-\alpha n}$, a shift $\boldsymbol{y} \in [-1,1]^n$ and a subspace $H \subseteq \mathbb{R}^n$ with $\dim(H) \geq \beta n$, finds an $\boldsymbol{x} \in (C \cdot K \cap H)$ with $\boldsymbol{x} + \boldsymbol{y} \in [-1,1]^n$ and $|\{i \in [n] : (\boldsymbol{x}+\boldsymbol{y})_i \in \{\pm 1\}\}| \geq (\beta - \gamma)n$.*

We give a proof in Section 1.3. We have the following corollary for full colorings. Here $K_S := K \cap \{x \in \mathbb{R}^n : x_i = 0, \forall i \notin S\}$.

Corollary 8. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex set. Given a function $f : [n] \to \mathbb{R}_{>0}$ with $\gamma_S(f(|S|) \cdot K_S) \geq 2^{-O(|S|)}$ for every $S \subseteq [n]$, there exists a randomized polynomial time algorithm to find a full coloring $x \in \{\pm 1\}^n$ so that $x \in \lambda K$, where $\lambda \lesssim \sum_{i=0}^{\lfloor \log n \rfloor} f(n/2^i)$. In particular, when $f(n) \lesssim n^\beta$ for some $\beta \leq 1$, we have $\lambda \lesssim \frac{1}{\beta} n^\beta$.*

Proof. Indeed, repeated iterations of Theorem 7 with $y_0 := 0$ and subsequent shifts y_{i+1} being the coordinates not reaching $\{-1, 1\}$ find $x := x_0 + \cdots + x_T \in \{\pm 1\}^n$ for $T := \lfloor \log n \rfloor$ with $x_t \in O(f(n/2^t)) \cdot K$. When $f(n) \lesssim n^\beta$, the summation is upper bounded by

$$\sum_{i=0}^{\infty} (n/2^i)^\beta = (1 - 2^{-\beta})^{-1} \cdot n^\beta \lesssim \frac{1}{\beta} \cdot n^\beta,$$

and this proves the statement. □

We also show that a weaker *hereditary* volume lower bound suffices to provide Gaussian measure lower bounds for arbitrary convex bodies. Previously such an implication was known only for the Gaussian measure of intersections with subspaces [53]:

Theorem 9. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Given $S \subseteq [n]$, denote by K_S the intersection with the coordinate subspace: $K_S := K \cap \{\boldsymbol{x} : x_i = 0 \; \forall i \notin S\} \subseteq \mathbb{R}^S$. Then*

$$\gamma_n(K) \geq \min_{S \subseteq [n]} \mathrm{vol}_{|S|}(K_S) \cdot 2^{-O(n)},$$

with the convention that $\mathrm{vol}_0(\{\mathbf{0}\}) = 1$. More generally, for any $\delta \in (0, 1]$,

$$\gamma_n(K) \geq \min_{S \subseteq [n], |S| \leq \delta n} \mathrm{vol}_{|S|}(K_S)^{1/\delta} \cdot 2^{-O(n/\delta)}.$$

We prove this in Section 1.4. In Chapter 6, we will need the following measure lower bound, which we prove in Section 1.5.

Theorem 10. *Suppose $K \subset \mathbb{R}^m$ is symmetric and convex, and that for a constant C_0,*

$$\gamma_m\left(\frac{C_0}{\alpha}K + \alpha\sqrt{m}B_2^m\right) \geq \frac{1}{2} \text{ for all } \alpha \in (0,1).$$

Then there is a constant C such that $\gamma_m(K) \geq \exp(-Cm)$.

We will also occasionally need to use vector balancing theorems of Banaszczyk related to Gaussian measure, which we cite without proof:

Theorem 11 (Banaszczyk [13]). *There exists a constant[1] $\beta > \frac{1}{5}$ so that for any convex body $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq \frac{1}{2}$ and vector $u \in \mathbb{R}^n$ with $\|u\|_2 \leq \beta$, there is a convex body $(K * u) \subseteq (K+u) \cup (K-u)$ with $\gamma_n(K*u) \geq \gamma_n(K)$.*

Theorem 12 (Banaszczyk [13]). *There is a constant $\alpha < 5$, so that for any $v_1, \ldots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$ for $i = 1, \ldots, T$ and any convex body $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq \frac{1}{2}$, there are signs $x_1, \ldots, x_T \in \{-1, 1\}$ so that*

$$\sum_{i=1}^{t} x_i v_i \in \alpha K.$$

Theorem 13 (Banaszczyk [14]). *There is a constant $\alpha < 5$, so that for any $v_1, \ldots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$ for $i = 1, \ldots, T$ and any convex body $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq 1 - \frac{1}{2T}$, there are signs $x_1, \ldots, x_T \in \{-1, 1\}$ so that*

$$\sum_{i=1}^{t} x_i v_i \in \alpha K \quad \forall t = 1, \ldots, T.$$

1.2 Facts about Gaussian measure

We also need one-dimensional Gaussian measure estimates, which we prove in Section 1.6.

[1]Banaszczyk's proof works as long as $\int_{-\beta}^{\beta} e^{-t^2/2} dt < \int_1^\infty e^{-t^2/2} dt$; for example take $\beta := 0.2001$.

Lemma 14. *For a strip $S = \{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{a}, \boldsymbol{x} \rangle| \leq 1\}$, one has*

$$\gamma_n(S) = \gamma_1(\{x \in \mathbb{R} : |x| \leq \|\boldsymbol{a}\|_2^{-1}\}) \geq 1 - \exp(-\|\boldsymbol{a}\|_2^{-2}/2).$$

Lemma 15. *For any $a \in \mathbb{R}^n$ with $\|a\|_2 \leq 1$ and $t \geq 1$ one has*

$$\Pr_{y \sim N(0, I_n)}[|\langle a, y \rangle| \leq t] \geq \exp(-e^{-t^2/2} \cdot \|a\|_2^2).$$

We use the following scaling lemma to deal with constant factors, see [163]:

Lemma 16. *Let $K \subset \mathbb{R}^n$ be a measurable set and B be a closed Euclidean ball such that $\gamma_n(K) = \gamma_n(B)$. Then $\gamma_n(tK) \geq \gamma_n(tB)$ for all $t \in [0, 1]$. In particular, if $\gamma_n(C \cdot K) \geq 2^{-O(n)}$ for some constant $C > 1$ then also $\gamma_n(K) \geq 2^{-O(n)}$.*

For teh next result, see [166].

Theorem 17. *If $F : \mathbb{R}^d \to \mathbb{R}$ is 1-Lipschitz, then for $t \geq 0$ one has*

$$\Pr_{\boldsymbol{y} \sim N(\boldsymbol{0}, \boldsymbol{I}_d)} \big[F(\boldsymbol{y}) > \mathbb{E}[F(\boldsymbol{y})] + t\big] \leq e^{-t^2/2}.$$

The classical *Urysohn Inequality* states that among all convex bodies of identical volume, the Euclidean ball minimizes the width:

Theorem 18 (Urysohn Inequality I)**.** *For any convex body $K \subseteq \mathbb{R}^n$ one has*

$$w(K) \geq 2 \cdot \Big(\frac{\mathrm{Vol}_n(K)}{\mathrm{Vol}_n(B_2^n)}\Big)^{1/n}.$$

We will need a variant that is phrased in terms of the Gaussian measure rather than volume. For a proof, see Eldan and Singh [61].

Theorem 19 (Gaussian Variant of Urysohn's Inequality)**.** *Let $K \subseteq \mathbb{R}^n$ be a convex body and let $r > 0$ be so that $\gamma_n(K) = \gamma_n(rB_2^n)$. Then $w(K) \geq w(rB_2^n) = r$.*

For a symmetric convex body K and a subspace H, the Gaussian measure of the section $K \cap (\boldsymbol{x} + H)$ is maximized when $\boldsymbol{x} = \boldsymbol{0}$ by log-concavity. Thus we have the following:

Lemma 20 (Gaussian measure of sections)**.** *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body and $H \subseteq \mathbb{R}^n$ a subspace. Then $\gamma_H(K \cap H) \geq \gamma_n(K)$.*

The following comparison inequality [95] will also be useful:

Lemma 21. *Let K be a symmetric convex body and let $0 \preceq A \preceq B$. Then*

$$\Pr_{y \sim N(0,A)}[y \in K] \geq \Pr_{y \sim N(0,B)}[y \in K].$$

1.3 Partial coloring via measure lower bound

In this section, we want to show the existence of partial fractional colorings for bodies K with $\gamma_n(K) \geq e^{-\alpha n}$ as promised in Theorem 7. The main innovation of this work compared to e.g. [146] is to handle an arbitrarily small constant $\alpha > 0$. We will show how to find a partial coloring that colors a small constant fraction of coordinates; then iterating the argument will color the promised $\beta - \gamma$ fraction. Also, instead of working with a shift $\boldsymbol{y}$ and a scaling of K, it will be notationally easier to work with a shifted and scaled box. Hence, for vectors $\boldsymbol{L}, \boldsymbol{R} \in \mathbb{R}^n_{\geq 0}$, we write $[-\boldsymbol{L}, \boldsymbol{R}] := [-L_1, R_1] \times \ldots \times [-L_n, R_n]$ as the box defined by constraints $-L_i \leq x_i \leq R_i$ for $i = 1, \ldots, n$. We use $N(\mathbf{0}, H)$ to denote the Gaussian distribution restricted to a subspace $H \subseteq \mathbb{R}^n$. Then the main technical result for this section will be:

Theorem 22. *For all constants $\alpha, \beta > 0$ there are $\varepsilon := \varepsilon(\alpha, \beta) > 0$ and $\delta := \delta(\alpha, \beta) > 0$ so that the following holds: Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body with $K \subseteq H$ for a subspace $H \subseteq \mathbb{R}^n$ with $\dim(H) \geq \beta n$ and $\gamma_H(K) \geq e^{-\alpha n}$; also let $\boldsymbol{L}, \boldsymbol{R} \in [0, \varepsilon]^n$. Assuming a weak separation oracle for K, there is a randomized polynomial time algorithm which finds an $\boldsymbol{x} \in K \cap [-\boldsymbol{L}, \boldsymbol{R}]$ so that $|\{i \in [n] : x_i \in \{-L_i, R_i\}\}| \geq \delta n$ with probability at least $1 - e^{-\Theta_{\varepsilon,\delta}(n)}$.*

Note that the considered box satisfies $[-\boldsymbol{L}, \boldsymbol{R}] \subseteq [-\varepsilon, \varepsilon]^n$. We would like to point out that applying the standard nonconstructive proof by Gluskin [68] and Giannopoulos [67] to a find a partial coloring $\boldsymbol{x} \in \{-\varepsilon, 0, \varepsilon\}^n$ with support $\Omega(n)$ will require either a *small* *enough* constant $\alpha > 0$, or ε needs to be exponentially small in n. In fact, it is not hard to construct a thin strip K with $\gamma_n(K) \geq e^{-\Omega(n)}$ so that K does not intersect $\{-1, 0, 1\}^n \setminus \{\mathbf{0}\}$ (even after a subexponential scaling). We show the construction in Section 1.7.

For our proof we make use of the mean width $w(Q) := \mathbb{E}_{\boldsymbol{\theta} \in S^{n-1}}[\sup_{\boldsymbol{x} \in Q} \langle \boldsymbol{\theta}, \boldsymbol{x} \rangle]$ of a body. We should point out that the connection between partial coloring arguments and mean width is due to Eldan and Singh [61]. Several of the claims require that n is chosen large enough.

Lemma 23. *Let $Q \subseteq \mathbb{R}^n$ be a symmetric convex body with $\gamma_n(Q) \geq e^{-\alpha n}$ for $\alpha > 0$. Then $w(Q) \geq \frac{1}{2} e^{-\alpha} \sqrt{n}$.*

Proof. Let $r > 0$ be the radius so that $\gamma_n(rB_2^n) = \gamma_n(Q)$. By *Urysohn's Inequality* (Theorem 19) one has $w(Q) \geq w(rB_2^n) = r$ so it suffices to give a lower bound on the radius r. A simple but useful estimate is that $2^n \leq \mathrm{Vol}_n(\sqrt{n} B_2^n) \leq 5^n$ for any $n \geq 1$. Moreover, the Gaussian density is maximized at $\gamma_n(\mathbf{0}) = \frac{1}{(\sqrt{2\pi})^n}$. Then for $\beta := 2e^{\alpha} \geq 2$ we have

$$\gamma_n\Big(\frac{\sqrt{n}}{\beta} B_2^n\Big) \leq \mathrm{Vol}_n\Big(\frac{\sqrt{n}}{\beta} B_2^n\Big) \cdot \gamma_n(\mathbf{0}) \leq \Big(\frac{5}{\beta}\Big)^n \cdot \frac{1}{(\sqrt{2\pi})^n} \leq \Big(\frac{2}{\beta}\Big)^n \stackrel{\beta = 2e^{\alpha}}{\leq} e^{-\alpha n}$$

and so $r \geq \frac{\sqrt{n}}{\beta} = \frac{\sqrt{n}}{2e^{\alpha}}$. $\square$

The key modification of our work in contrast to [146] is a finer upper bound on the distance of a Gaussian to K:

Lemma 24. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex set with $\gamma_n(K) \geq e^{-\alpha n}$ where $\alpha \geq 1$ and n is large enough. Then*

$$\mathop{\mathbb{E}}_{\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{I}_n)}[d(\boldsymbol{x}, K)] \leq \sqrt{n} \cdot \Big(1 - \frac{1}{512 \alpha e^{4\alpha}}\Big)$$

Proof. Note that by Theorem 17 we have $\Pr_{\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{I}_n)}[\|\boldsymbol{x}\|_2 \geq 4\sqrt{\alpha n}] \leq e^{-2\alpha n}$, hence the restriction $Q := K \cap 4\sqrt{\alpha n} B_2^n$ still has $\gamma_n(Q) \geq \gamma_n(K) - e^{-2\alpha n} \geq e^{-2\alpha n}$ for n large enough. Then by the previous Lemma we know that $w(Q) \geq \frac{\sqrt{n}}{2e^{2\alpha}}$. For a vector $\boldsymbol{x}$, let $\boldsymbol{z}(\boldsymbol{x}) := \mathrm{argmax}\{\langle \boldsymbol{z}, \boldsymbol{x} \rangle : \boldsymbol{z} \in Q\}$. As we just showed, $\mathbb{E}_{\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{I}_n)}[\Big\langle \boldsymbol{z}(\boldsymbol{x}), \frac{\boldsymbol{x}}{\|\boldsymbol{x}\|_2} \Big\rangle] \geq \frac{\sqrt{n}}{2e^{2\alpha}}$. Let $\lambda \in [0, 1]$ be a parameter that we determine later. Note that the point $\lambda \cdot \boldsymbol{z}(\boldsymbol{x})$ lies in Q.

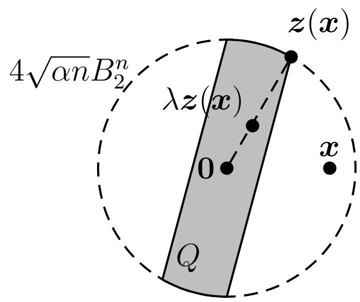

This point can be used to bound

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{x}\sim N(\mathbf{0},\boldsymbol{I}_n)}[\|\boldsymbol{x}-\lambda \boldsymbol{z}(\boldsymbol{x})\|_2^2] &= \mathbb{E}[\|\boldsymbol{x}\|_2^2] - 2\lambda\,\mathbb{E}[\langle \boldsymbol{x},\boldsymbol{z}\rangle] + \mathbb{E}[\lambda^2\|\boldsymbol{z}\|_2^2] \\
&= \underbrace{\mathbb{E}[\|\boldsymbol{x}\|_2^2]}_{=n} - 2\lambda \underbrace{\mathbb{E}[\|\boldsymbol{x}\|_2]}_{\geq \frac{1}{2}\sqrt{n}} \cdot \underbrace{\mathop{\mathbb{E}}_{\boldsymbol{\theta}\in S^{n-1}}[\langle \boldsymbol{\theta},\boldsymbol{z}(\boldsymbol{\theta})\rangle]}_{\geq \sqrt{n}/(2e^{2\alpha})} + \mathbb{E}[\lambda^2 \underbrace{\|\boldsymbol{z}\|_2^2}_{\leq 16\alpha n}] \\
&\leq n - \frac{1}{2}e^{-2\alpha}\lambda n + \lambda^2\cdot 16\alpha n \stackrel{\lambda:=\frac{1}{64\alpha e^{2\alpha}}}{=} n\cdot\Big(1-\frac{1}{256\alpha e^{4\alpha}}\Big)
\end{aligned}$$

Then

$$\mathbb{E}[d(\boldsymbol{x},Q)] \stackrel{\lambda\boldsymbol{z}\in Q}{\leq} \mathbb{E}[\|\boldsymbol{x}-\lambda\boldsymbol{z}\|_2] \stackrel{\text{Jensen}}{\leq} \mathbb{E}[\|\boldsymbol{x}-\lambda\boldsymbol{z}\|_2^2]^{1/2} \leq \sqrt{n}\cdot\sqrt{1-\frac{1}{256\alpha e^{4\alpha}}} \leq \sqrt{n}\cdot\Big(1-\frac{1}{512\alpha e^{4\alpha}}\Big)$$

using $\sqrt{1-y} \leq 1-\frac{y}{2}$ for $0 \leq y \leq 1$. $\square$

Lemma 24 can be extended to the case that K is included in a not too small subspace H.

Lemma 25. *Let $\alpha \geq 1$, $0 < \beta \leq 1$ be constants. Let $H \subseteq \mathbb{R}^n$ be a subspace with $\dim(H) \geq \beta n$ and let $K \subseteq H$ be a symmetric convex body with $\gamma_H(K) \geq e^{-\alpha n}$. For n large enough, one has*

$$\mathop{\mathbb{E}}_{\boldsymbol{x}\sim N(\mathbf{0},\boldsymbol{I}_n)}[d(\boldsymbol{x},K)] \leq \sqrt{n}\cdot\Big(1-\frac{\beta}{512\alpha e^{4\alpha}}\Big)$$

Proof. Note that one can generate a Gaussian $\boldsymbol{x} \sim N(\mathbf{0},\boldsymbol{I}_n)$ as $\boldsymbol{x} = \boldsymbol{x}_1 + \boldsymbol{x}_2$ where $\boldsymbol{x}_1 \sim N(\mathbf{0},H^\perp)$ and $\boldsymbol{x}_2 \sim N(\mathbf{0},H)$ independently. Then $d(\boldsymbol{x},K)^2 = d(\boldsymbol{x}_1,H)^2 + d(\boldsymbol{x}_2,K)^2$ by

Pythagoras. Hence

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{x}\sim N(\mathbf{0},\boldsymbol{I}_n)}\left[d(\boldsymbol{x},K)^2\right] &\leq \mathop{\mathbb{E}}_{\boldsymbol{x}_1\sim N(\mathbf{0},H^{\perp})}\left[d(\boldsymbol{x}_1,H)^2\right] + \mathop{\mathbb{E}}_{\boldsymbol{x}_2\sim N(\mathbf{0},H)}\left[d(\boldsymbol{x}_2,K)^2\right] \\
&\stackrel{\text{Lem 24}}{\leq} \dim(H^{\perp}) + \dim(H)\cdot\Big(1-\frac{1}{256\alpha e^{4\alpha}}\Big) \\
&\stackrel{\dim(H)\geq\beta n}{\leq} n\cdot\Big(1-\frac{\beta}{256\alpha e^{4\alpha}}\Big)
\end{aligned}$$

As in the proof of Lemma 24, the claim follows after applying Jensen inequality with the fact that $\sqrt{1-y}\leq 1-\frac{y}{2}$ for $0\leq y\leq 1$. □

Next, we show the average distance of a Gaussian to the cube $[-\varepsilon,\varepsilon]^n$ is $\sqrt{n}\cdot(1-\Theta(\varepsilon))$.

Lemma 26. *Let $\varepsilon>0$. Then for n large enough one has*

$$\Pr_{\boldsymbol{x}\sim N(\mathbf{0},\boldsymbol{I}_n)}\Big[d(\boldsymbol{x},[-\varepsilon,\varepsilon]^n)\geq(1-5\varepsilon)\sqrt{n}\Big]\geq 1-\exp\Big(-\frac{\varepsilon^2}{2}n\Big)$$

Proof. Let $\boldsymbol{y}:=\boldsymbol{y}(\boldsymbol{x}):=\text{argmin}\{\|\boldsymbol{x}-\boldsymbol{y}\|_2:\boldsymbol{y}\in[-\varepsilon,\varepsilon]^n\}$ be the closest point in the cube to $\boldsymbol{x}$. For an individual coordinate $i\in[n]$ the expected contribution to the distance is

$$\mathbb{E}\left[d(x_i,[-\varepsilon,\varepsilon])^2\right]=\mathbb{E}\left[|x_i-y_i|^2\right]=\underbrace{\mathbb{E}[x_i^2]}_{=1}-2\underbrace{\mathbb{E}[x_iy_i]}_{\leq\varepsilon\,\mathbb{E}[|x_i|]}+\underbrace{\mathbb{E}[y_i^2]}_{\geq 0}\geq 1-2\sqrt{\frac{2}{\pi}}\cdot\varepsilon\geq 1-2\varepsilon.$$

Then by linearity $\mathbb{E}[d(\boldsymbol{x},[-\varepsilon,\varepsilon]^n)^2]^{1/2}\geq\sqrt{n\cdot(1-2\varepsilon)}\geq\sqrt{n}\cdot(1-2\varepsilon)$. Recall that the distance function $F(\boldsymbol{x}):=d(\boldsymbol{x},[-\varepsilon,\varepsilon]^n)$ is 1-Lipschitz and for such functions the difference $|\,\mathbb{E}[F(\boldsymbol{x})]-\mathbb{E}[F(\boldsymbol{x})^2]^{1/2}|$ is bounded by an absolute constant. Then $\mathbb{E}[F(\boldsymbol{x})]\geq\sqrt{n}\cdot(1-4\varepsilon)$ for n large enough. Finally by Theorem 17 one has $\Pr[F(\boldsymbol{x})<\mathbb{E}[F(\boldsymbol{x})]-\varepsilon\sqrt{n}]\leq e^{-\varepsilon^2n/2}$ for $\boldsymbol{x}\sim N(\mathbf{0},\boldsymbol{I}_n)$ which then gives the claim as $\mathbb{E}[F(\boldsymbol{x})]-\varepsilon\sqrt{n}\geq(1-5\varepsilon)\sqrt{n}$. □

We will now prove Theorem 22. Let $H\subseteq\mathbb{R}^n$ be a subspace with $\dim(H)\geq\beta n$ and let $K\subseteq H\subseteq\mathbb{R}^n$ be a symmetric convex body with $\gamma_H(K)\geq e^{-\alpha n}$. Moreover, let $L_i,R_i\in[0,\varepsilon]$ be given parameters where the choice of $\varepsilon:=\varepsilon(\alpha,\beta)>0$ will be made in the upcoming proof of Lemma 27. We will use the following algorithm:

(1) Pick $\boldsymbol{x}^* \sim N(\mathbf{0}, \boldsymbol{I}_n)$ at random.

(2) Compute $\boldsymbol{y}^* := \text{argmin}\big\{\|\boldsymbol{x}^* - \boldsymbol{y}\|_2 : \boldsymbol{y} \in K \cap [-\boldsymbol{L}, \boldsymbol{R}]^n\big\}$.

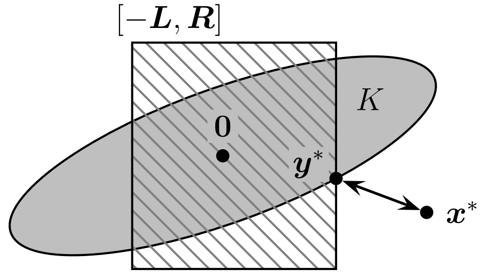

Note that the step (2) is a convex program which can be solved in polynomial time, see [70]. Now we can finish the proof of Theorem 22.

Lemma 27. *If $\varepsilon, \delta > 0$ are chosen small enough (depending on α), then with probability $1 - e^{-\Omega_{\varepsilon,\delta}(n)}$ one has $|\{i \in [n] : y_i^* \in \{-L_i, R_i\}\}| \geq \delta n$.*

Proof. For a set of indices $I \subseteq [n]$ we abbreviate the subspace $H(I) := \{\boldsymbol{x} \in H \mid x_i = 0\ \forall i \in I\}$. Moreover we abbreviate $K(I) := \{\boldsymbol{x} \in K \mid -L_i \leq x_i \leq R_i\ \forall i \in I\}$ as the intersection of K with the slabs corresponding to coordinates in I. Consider the two events

$$\mathcal{E}_1 := \text{“}d(\boldsymbol{x}^*, K \cap [-\boldsymbol{L}, \boldsymbol{R}]) \geq (1 - 5\varepsilon) \cdot \sqrt{n}\text{”}$$
$$\mathcal{E}_2 := \text{“for all } I \subseteq [n] \text{ with } |I| \leq \delta n \text{ one has } d(\boldsymbol{x}^*, K \cap H(I)) \leq (1 - 10\varepsilon)\sqrt{n}\text{”}$$

We will see that both events $\mathcal{E}_1$ and $\mathcal{E}_2$ happen with overwhelming probability.

Claim I. *One has* $\Pr[\mathcal{E}_1] \geq 1 - \exp(-\frac{\varepsilon^2}{2}n)$.

Proof of Claim I. Follows from Lemma 26 as $d(\boldsymbol{x}^*, K \cap [-\boldsymbol{L}, \boldsymbol{R}]) \geq d(\boldsymbol{x}^*, K \cap [-\varepsilon, \varepsilon]^n) \geq d(\boldsymbol{x}^*, [-\varepsilon, \varepsilon]^n)$.

Claim II. *If $\varepsilon, \delta > 0$ are small enough, then* $\Pr[\mathcal{E}_2] \geq 1 - e^{-\Theta_\varepsilon(n)}$.

Proof of Claim II. For any index set I one can lower bound the measure as $\gamma_{H(I)}(K \cap H(I)) \geq \gamma_H(K) \geq e^{-\alpha n}$ by Lemma 20. Let us abbreviate $\mathcal{I} := \{I \subseteq [n] : |I| \leq \delta n\}$ as the family of small index sets. For $I \in \mathcal{I}$ we have $\dim(H(I)) \geq \dim(H) - |I| \geq \frac{\beta}{2}n$, if we choose $\delta \leq \frac{\beta}{2}$. Then by Lemma 25 we know that a fixed $I \in \mathcal{I}$ has $\mathbb{E}_{\boldsymbol{x} \sim N(\mathbf{0}, \boldsymbol{I}_n)}[d(\boldsymbol{x}, K \cap H(I))] \leq$

$\sqrt{n} \cdot \big(1 - \frac{\beta/2}{512 \cdot \alpha e^{4\alpha}}\big) \leq (1-20\varepsilon)\sqrt{n}$, if we choose $\varepsilon \leq \frac{\beta/2}{20 \cdot 512 \alpha e^{4\alpha}}$. Then by concentration one has $\Pr_{\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)}[d(\boldsymbol{x}, K \cap H(I)) > (1-10\varepsilon)\sqrt{n}] \leq \exp(-50\varepsilon^2 n)$, see Theorem 17. A useful bound is $|\mathcal{I}| \leq e^{2\delta \log_2(\frac{1}{\delta})n} \leq e^{\varepsilon^2 n}$ if we choose δ small enough compared to ε. Then

$$\begin{aligned} \Pr[\mathcal{E}_2] &\overset{\text{union bound}}{\leq} \sum_{I \in \mathcal{I}} \Pr\left[d(\boldsymbol{x}^*, K \cap H(I)) > (1-10\varepsilon)\sqrt{n}\right] \\ &\leq e^{\varepsilon^2 n} \cdot \exp(-50\varepsilon^2 n) \leq \exp\Big(-40\varepsilon^2 n\Big). \quad \square \end{aligned}$$

Now we have everything to finish the proof. Fix an outcome of the vector $\boldsymbol{x}^*$ so that the events $\mathcal{E}_1$ and $\mathcal{E}_2$ are both true, and abbreviate $I^* := \{i \in [n] : y_i^* \in \{-L_i, R_i\}\}$. Suppose for the sake of contradiction that $|I^*| < \delta n$. Then

$$\begin{aligned} (1-10\varepsilon)\sqrt{n} &\overset{\mathcal{E}_2 \text{ true } \& \ I^* \in \mathcal{I}}{\geq} d(\boldsymbol{x}^*, K \cap H(I^*)) \\ &\overset{K \cap H(I^*) \subseteq K(I^*)}{\geq} d(\boldsymbol{x}^*, K(I^*)) \\ &\overset{(*)}{=} d(\boldsymbol{x}^*, K \cap [-\boldsymbol{L}, \boldsymbol{R}]) \\ &\overset{\mathcal{E}_1 \text{ true}}{\geq} (1-5\varepsilon)\sqrt{n} \end{aligned}$$

which is a contradiction. Here the crucial argument for $(*)$ is that $d(\boldsymbol{x}^*, K \cap [-\boldsymbol{L}, \boldsymbol{R}]) = \min\{\|\boldsymbol{x}^* - \boldsymbol{y}\|_2 : \boldsymbol{y} \in K \text{ and } -L_i \leq y_i \leq R_i \ \forall i \in [n]\}$ is a *convex minimization* problem and the optimum value will not change if linear constraints are discarded that are not tight for the optimum $\boldsymbol{y}^*$, and the box constraints for coordinates $I^* \setminus [n]$ are indeed not tight. $\square$

We stated such a result earlier in Theorem 7. Now we are ready to prove it:

Proof of Theorem 7. The basic idea is to simply apply Theorem 22 a constant number of times until the desired number of elements is colored. We assume $\beta > \gamma$ since otherwise there is nothing to prove. Let $\varepsilon := \varepsilon(\alpha, \gamma), \delta := \delta(\varepsilon, \gamma) > 0$ be the constants from Theorem 22 that work for the given α and $\beta' := \gamma > 0$.

We set $\boldsymbol{y}^{(0)} := \boldsymbol{y}$ and for $t \geq 0$ we set $F^{(t)} := \{i \in [n] : y_i^{(t)} \in \{-1, 1\}\}$ as the variables that are *frozen*. Suppose for some t we have constructed a sequence $\boldsymbol{y}^{(0)}, \ldots, \boldsymbol{y}^{(t)}$ and still $|F^{(t)}| < (\beta - \gamma)n$. Set $H^{(t)} := \{\boldsymbol{x} \in H \mid x_i = 0 \ \forall i \in F^{(t)}\}$ be the subspace of H where we fix frozen coordinates to be 0. Note that $\dim(H^{(t)}) \geq \dim(H) - |F^{(t)}| \geq \gamma n$. Moreover $\gamma_{H^{(t)}}(K \cap H^{(t)}) \geq \gamma_H(K) \geq e^{-\alpha n}$ by Lemma 20. We set $R_i := \frac{\varepsilon}{2} \cdot (1 - y_i^{(t)})$ and $L_i := \frac{\varepsilon}{2} \cdot (y_i^{(t)} - (-1))$ for $i \in [n] \setminus F^{(t)}$ and $R_i := L_i := \varepsilon$ for $i \in F^{(t)}$ and apply Theorem 22. With high probability, the algorithm succeeds and provides a vector $\boldsymbol{x}^{(t)}$. We update $\boldsymbol{y}^{(t+1)} := \boldsymbol{y}^{(t)} + \frac{2}{\varepsilon}\boldsymbol{x}^{(t)} \in [-1,1]^n$ where $\|\boldsymbol{y}^{(t+1)}\|_K \leq \|\boldsymbol{y}^{(t)}\|_K + \frac{2}{\varepsilon}$ by the triangle inequality. Moreover, the number of frozen coordinates increases[2] to $|F^{(t+1)}| \geq |F^{(t)}| + \delta n$. We will terminate after at most $\frac{1}{\delta}$ iterations and if T is the final iteration, then $\boldsymbol{y}^{(T)} \in [-1,1]^n \cap \frac{2}{\varepsilon\delta} K$ as desired. $\square$

We would like to mention that Theorem 7 may also be deduced, after some work, from the Gaussian measure amplification techniques derived in [53] with the use of α-regular M-ellipsoids. We believe the analysis presented here is simpler, since the existence of such regular M-ellipsoids is a deep result in convex geometry.

1.4 From hereditary volume bounds to Gaussian measure

This section is devoted to the proof of Theorem 9, which provides a connection between hereditary volume and Gaussian measure. For a brief motivation, note that for any convex body $K \subseteq \mathbb{R}^n$ and any $S \subseteq [n]$ one has $\mathrm{vol}_{|S|}(K_S) \geq \gamma_{|S|}(K_S) \geq \gamma_n(K)$. It is therefore a natural question whether a converse holds, and Theorem 9 shows that this is indeed the case. As a corollary, we settle up to an exponential factor a conjecture of [30] that coordinate sections minimize the Gaussian measure among all sections of scaled ℓ_p balls.

We would also like to mention that we cannot hope for a refinement of the right side to only sections of dimension δn. For example when $K = \varepsilon \cdot B_2^{\delta n - 1} \times \mathbb{R}^{n - \delta n + 1}$, all δn-dimensional sections have infinite volume yet $\gamma(K) \to 0$ as $\varepsilon \to 0$.

While relatively short, our proof does use several auxilliary results. The key ingredient is the following formula which expresses the volume of the Minkowski sum of a convex body and an Euclidean ball as a weighted sum of *quermassintegrals* $W_i(K)$ which are average volumes of projections. Recall that given $A, B \subseteq \mathbb{R}^n$, $A + B := \{\boldsymbol{a} + \boldsymbol{b} : \boldsymbol{a} \in A, \boldsymbol{b} \in B\}$.

[2] For frozen coordinates i we did set $L_i = R_i = \varepsilon$ so that $\boldsymbol{x}^{(t)}$ will indeed contain δn "fresh" coordinates that become tight, rather than rediscovering the coordinates in $F^{(t)}$.

Lemma 28 (Kubota's Integral Formula [133])**.** *For any convex body $K \subset \mathbb{R}^n$, we have*

$$\mathrm{vol}_n(K + \lambda B_2^n) = \sum_{i=0}^{n} \lambda^i \binom{n}{i} W_i(K)$$

with

$$W_{n-i}(K) := \frac{\mathrm{vol}_n(B_2^n)}{\mathrm{vol}_i(B_2^i)} \int_{G(n,i)} \mathrm{vol}_i(\pi_L(K))\, \mathrm{d}L,$$

where the integral is over the uniform measure over $G(n,i)$, which is the set of i-dimensional linear subspaces $L \subseteq \mathbb{R}^n$ and $\pi_L(K)$ denotes the orthogonal projection of K onto L.

In order to relate projections to slices, we use polarity. Given a symmetric convex set $K \subseteq \mathbb{R}^n$, its *polar* is $K^\circ := \{\boldsymbol{y} \in \mathrm{span}(K) \mid \langle \boldsymbol{x}, \boldsymbol{y} \rangle \leq 1 \ \forall \boldsymbol{x} \in K\}$. The following lemma elucidates the reason polars are helpful to transform projections into slices:

Lemma 29. *Given a symmetric convex body $K \subseteq \mathbb{R}^n$ and any subspace $H \subseteq \mathbb{R}^n$, we have $(K \cap H)^\circ = \pi_H(K^\circ)$.*

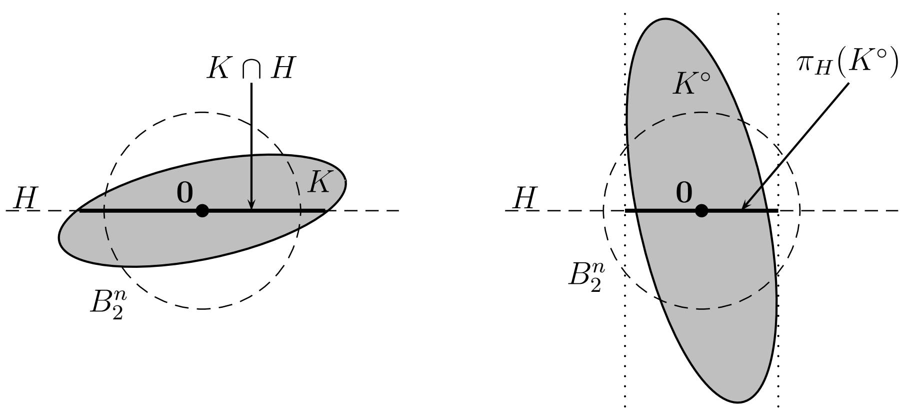

It is also well-known that polarity transforms intersections into convex hulls:

Lemma 30. *Given symmetric convex bodies $K, L \subseteq \mathbb{R}^n$, we have $(K \cap L)^\circ = \mathrm{conv}(K^\circ, L^\circ)$.*

For a detailed introduction to polarity we refer to Rockefellar [142]. Finally, we need the Blaschke-Santaló Inequality and its deep converse due to Bourgain-Milman [7]:

Lemma 31. *Given a symmetric convex body $K \subseteq \mathbb{R}^n$, we have $2^{O(n)} \geq \frac{\mathrm{vol}_n(K) \cdot \mathrm{vol}_n(K^\circ)}{\mathrm{vol}_n(B_2^n)^2} \geq 2^{-O(n)}$.*

The starting point of the proof, which connects the Gaussian measure to the Minkowski sum with an Euclidean ball, is given by the following bound:

Lemma 32. *Given a symmetric convex body $K \subseteq \mathbb{R}^n$, $\gamma_n(K) \geq \mathrm{vol}_n\Big(K^\circ + \frac{1}{\sqrt{n}}B_2^n\Big)^{-1} \cdot n^{-n} \cdot 2^{O(n)}$.*

Proof. We start by noting that we can lower bound the Gaussian measure upon restriction to a $\sqrt{n}$-radius ball:

$$\gamma_n(K) = \frac{1}{(2\pi)^{n/2}} \int_K e^{-\|\boldsymbol{x}\|_2^2/2}\, \mathrm{d}\boldsymbol{x} \geq \frac{1}{(2\pi e)^{n/2}} \mathrm{vol}_n(K \cap \sqrt{n}B_2^n),$$

and since $(K \cap \sqrt{n}B_2^n)^\circ = \mathrm{conv}(K^\circ, \frac{1}{\sqrt{n}}B_2^n)$ by Lemma 30, we conclude

$$\begin{aligned} \gamma_n(K) &\geq \mathrm{vol}_n(K \cap \sqrt{n}B_2^n) \cdot 2^{-O(n)} \\ &\overset{\text{Lem 31}}{\geq} \mathrm{vol}_n\Big(\mathrm{conv}\Big(K^\circ, \frac{1}{\sqrt{n}}B_2^n\Big)\Big)^{-1} \cdot n^{-n} \cdot 2^{-O(n)} \\ &\geq \mathrm{vol}_n\Big(K^\circ + \frac{1}{\sqrt{n}}B_2^n\Big)^{-1} \cdot n^{-n} \cdot 2^{O(n)}, \end{aligned}$$

since $\mathrm{conv}(K^\circ, \frac{1}{\sqrt{n}}B_2^n) \subseteq K^\circ + \frac{1}{\sqrt{n}}B_2^n$. ◻

In order to connect slices to *coordinate* slices, we apply a result of [53] for ellipsoids. Thus we will need to use the existence of M-ellipsoids [7]:

Lemma 33. *For any symmetric convex body $K \subseteq \mathbb{R}^n$ there exists an ellipsoid $E \subseteq \mathbb{R}^n$ for which there exist collections of centers $S_E, S_K \subseteq \mathbb{R}^n$ with $|S_E|, |S_K| \leq 2^{O(n)}$ so that $K \subseteq \bigcup_{c \in S_E}(c + E)$ and $E \subseteq \bigcup_{c' \in S_K}(c' + K)$.*

Proof of the first inequality in Theorem 9. Kubota's integral formula (Lemma 28) applied to K° yields

$$W_{n-i}(K^\circ) = \frac{\mathrm{vol}_n(B_2^n)}{\mathrm{vol}_i(B_2^i)} \int_{G(n,i)} \mathrm{vol}_i(\pi_L(K^\circ))\, \mathrm{d}L.$$

By Lemma 29 and Santaló's inequality (Lemma 31) we know that for any subspace L,

$$\mathrm{vol}_i(\pi_L(K^\circ)) \leq \mathrm{vol}_i(B_2^i)^2 \cdot \mathrm{vol}_i(K \cap L)^{-1} \leq M^{-1} \cdot i^{-i} \cdot 2^{O(i)},$$

where we choose to denote $M := \min_{\dim L = i \leq n} \mathrm{vol}_i(K \cap L)$. We conclude

$$W_{n-i}(K^\circ) \le M^{-1} \cdot n^{-n/2} \cdot i^{-i/2} \cdot 2^{O(n)},$$

so that

$$n^{-(n-i)/2} \cdot W_{n-i}(K^\circ) \le M^{-1} \cdot n^{-n} \cdot 2^{O(n)},$$

by using $(n/i)^i \le 2^{O(n)}$ for $i \in [n]$. Taking $\lambda := 1/\sqrt{n}$ and summing over $i \in [n]$ in Lemma 28 gives

$$\mathrm{vol}_n\Big(K^\circ + \frac{1}{\sqrt{n}} B_2^n\Big) \le M^{-1} \cdot n^{-n} \cdot 2^{O(n)},$$

so that by Lemma 32 we obtain $\gamma_n(K) \ge M \cdot 2^{-O(n)}$. It remains to show that the minimal *coordinate* sections are not much larger than the minimal sections. With this purpose in mind, let E be an M-ellipsoid of K. By Lemma 33, there exist collections S_E, S_K with $|S_E|, |S_K| \le 2^{O(n)}$ so that $K \subseteq \bigcup_{c \in S_E}(c + E)$ and $E \subseteq \bigcup_{c' \in S_K}(c' + K)$. Note that for any i-dimensional subspace L we have

$$\mathrm{vol}_i(K \cap L) \le \sum_{c \in S_E} \mathrm{vol}_i((c + E) \cap L) \le 2^{O(n)} \cdot \mathrm{vol}_i(E \cap L)$$

and similarly

$$\mathrm{vol}_i(E \cap L) \le \sum_{c' \in S_K} \mathrm{vol}_i((c' + K) \cap L) \le 2^{O(n)} \cdot \mathrm{vol}_i(K \cap L),$$

where by Brunn's concavity principle the sections with largest volume are those through the origin. Thus it suffices to show that

$$\min_{\dim L = i} \mathrm{vol}_i(E \cap L) \ge \min_{S \subseteq [n], |S| = i} \mathrm{vol}_i(E_S) \cdot 2^{-O(n)}.$$

Indeed this follows a form of restricted invertibility in the work of Dadush, Nikolov, Talwar and Tomczak-Jaegermann, who showed in [53] (see p. 8) an improved bound of

$$\min_{\dim L = i} \mathrm{vol}_i(E \cap L) \ge \min_{S \subseteq [n], |S| = i} \mathrm{vol}_i(E_S) \cdot \binom{n}{i}^{-1}. \qquad \square$$

We now prove the second part of Theorem 9 which restricts our attention to sections of dimension $\leq \delta n$. For this we need the following inequality for quermassintegrals which can be seen as a strenghtening of the isoperimetric inequality:

Theorem 34 (Alexandrov Inequality [133])**.** *Given $i \geq j$ we have*

$$\Big(\frac{W_{n-i}(K)}{\mathrm{vol}_i(B_2^i)}\Big)^{1/i} \leq \Big(\frac{W_{n-j}(K)}{\mathrm{vol}_j(B_2^j)}\Big)^{1/j}.$$

Proof of the second inequality in Theorem 9. We proceed as in the proof of the first inequality. Setting $\lambda := 1/\sqrt{n}$ we still have, for $j \leq \delta n$,

$$\begin{aligned}\lambda^{n-j} W_{n-j}(K^\circ) &\leq \max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1}(K \cap L) \cdot n^{-n} \cdot 2^{O(n)} \\ &\leq \max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1/\delta}(K \cap L) \cdot n^{-n} \cdot 2^{O(n)},\end{aligned}$$

as the maximum is at least one (for $i = 0$). For $j > \delta n$ we use Theorem 34 to see that

$$\lambda^{n-j} W_{n-j}(K^\circ) \leq \lambda^{n-j} (W_{n-\delta n}(K^\circ))^{j/(\delta n)} \cdot \mathrm{vol}_j(B_2^j) \cdot \mathrm{vol}_{\delta n}(B_2^{\delta n})^{-j/\delta n}$$

and proceed as in the first half of the proof:

$$\begin{aligned}\lambda^{n-j} W_{n-j}(K^\circ) &\leq \lambda^{n-j} \cdot (W_{n-\delta n}(K^\circ))^{j/(\delta n)} \cdot \mathrm{vol}_j(B_2^j) \cdot \mathrm{vol}_{\delta n}(B_2^{\delta n})^{-j/\delta n} \\ &\leq \lambda^{n-j} \cdot \Big(\max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1}(K \cap L) \cdot n^{-n/2} \cdot (\delta n)^{-\delta n/2}\Big)^{j/(\delta n)} \cdot (\delta n/j)^{j/2} \cdot 2^{O(n/\delta)} \\ &\leq \lambda^{n-j} \cdot n^{-j/(2\delta)} \cdot j^{-j/2} \cdot \max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1/\delta}(K \cap L) \cdot 2^{O(n/\delta)} \\ &= n^{-n/2} \cdot n^{-j/(2\delta)} \cdot \underbrace{(n/j)^{j/2}}_{\leq 2^{O(n)}} \max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1/\delta}(K \cap L) \cdot 2^{O(n/\delta)} \\ &\leq n^{-n} \cdot \max_{\dim L = i \leq \delta n} \mathrm{vol}_i^{-1/\delta}(K \cap L) \cdot 2^{O(n/\delta)}.\end{aligned}$$

The statement follows as before: by summing over $j \in [n]$ in Lemma 28 we obtain

$$\gamma_n(K) \geq \min_{\dim L = i \leq \delta n} \mathrm{vol}_i^{1/\delta}(K \cap L) \cdot 2^{-O(n/\delta)},$$

and we can pass to coordinate sections via M-ellipsoids. $\square$

Remark 1. *Barthe, Guédon, Mendelson, and Naor conjectured that coordinate slices maximize the Gaussian volume among all slices of a (scaled) ℓ_p ball [30] (see the remark in p. 28). We can use the above result to give an affirmative answer up to $2^{-O(n)}$:*

Corollary 35. *Let $p \geq 2$, $r > 0$ and $H \subseteq \mathbb{R}^d$ an n-dimensional subspace. Then*

$$\gamma_H(rB_p^d \cap H) \geq \gamma_n(rB_p^n) \cdot 2^{-O(n)}.$$

Proof. If $r > n^{1/p}$, the right side is already $2^{-O(n)}$ so we may assume that $r \leq n^{1/p}$. A well-known result of Meyer-Pajor asserts that coordinate sections minimize the *volume* among all sections of the ℓ_p ball [114]. Applying Theorem 9 and using Meyer-Pajor we get

$$\gamma_H(rB_p^d \cap H) \geq \min_{L \subseteq H, \dim L = i} \mathrm{vol}_i(rB_p^d \cap L) \geq \min_{i \leq n} \mathrm{vol}_i(rB_p^i) \geq \gamma_n(rB_p^n) \cdot 2^{-O(n)}. \qquad \square$$

Remark 2. *We mention another application of Theorem 9. For a symmetric convex $K \subseteq \mathbb{R}^n$, denote the hereditary discrepancy* $\mathrm{hd}(K)$ *as the minimum $t \geq 0$ so that tK_S intersects $\{-1,1\}^S \times \{0\}^{[n]\setminus S}$ for all $S \subseteq [n]$. In [53] it is shown that we have a lower bound* $\mathrm{hd}(K) \geq \max_{S \subseteq [n]} \inf\{t : \mathrm{vol}_{|S|}(tK_S) \geq 1\}$*, where the left side is known as the* volume lower bound $\mathrm{volLB}(K)$*. In fact an analogous argument also shows the lower bound* $\mathrm{hd}(K) \geq \max_{S \subseteq [n]} \inf\{t : \gamma_{|S|}(tK_S) \geq 2^{-C|S|}\}$ *for a universal constant $C > 0$. Since the volume of a convex body is always lower bounded by its Gaussian measure, this lower bound is at least* $\mathrm{volLB}(K)$ *up to a factor of 2^C. Theorem 9 immediately implies that it is also at most* $\mathrm{volLB}(K)$ *up to a constant.*

1.5 Gaussian measure lower bounds from expansion

In this section, we show Theorem 10. We will need several facts about covering numbers. We start with the definition:

Definition 36 (Covering Numbers)**.** *For two convex bodies $K, T \subseteq \mathbb{R}^n$, we define the covering number $\mathcal{N}(K,T)$ as the minimum number N such that there exist centers $x_1, \ldots, x_N \in \mathbb{R}^n$ with $K \subseteq \cup_{i=1}^N (x_i + T)$, i.e. K can be covered by N translates of T.*

We need the following few standard facts about covering numbers (see [7]).

Lemma 37 (Volume Bounds for Covering Numbers)**.** *Given convex bodies $K, T \subseteq \mathbb{R}^n$. If T is symmetric, we have $\frac{\mathrm{vol}_n(K)}{\mathrm{vol}_n(T)} \leq \mathcal{N}(K, T) \leq 2^n \cdot \frac{\mathrm{vol}_n(2K+T)}{\mathrm{vol}_n(T)}$.*

Lemma 38 (Symmetrization)**.** *Let $K \subseteq \mathbb{R}^n$ be a convex body. Then $\mathcal{N}(K - K, K) \leq 2^{O(n)}$.*

The main technical tool we use is the existence of 1*-regular M-ellipsoids*. We state a variant of [133], which gives a result for p-regular M-ellipsoids for all $p \in (0, 2)$; we only use $p = 1$.

Proposition 39 (Corollary 7.16, [133])**.** *There is a constant C_1 such that for any symmetric convex $K \subseteq \mathbb{R}^m$, there is an ellipsoid $\mathcal{E}$ so that for all $t > 0$,*

$$\max\{N(K, t\mathcal{E}), N(\mathcal{E}, tK)\} \leq \exp(C_1 m t^{-1}).$$

For the following standard facts, see Theorem 4.1.13 and Facts 4.1.7, 4.1.8 and 4.1.9 in [7].

Fact 40. *For any convex sets $A, B, C \subseteq \mathbb{R}^m$, we have $N(A, B) \leq N(A, C)N(C, B)$.*

Fact 41. *For any convex sets $A, B, C \subseteq \mathbb{R}^m$, we have $N(A + C, B + C) \leq N(A, B)$.*

Fact 42. *For any convex sets $A, B \subseteq \mathbb{R}^m$, we have $N(A, B) \geq \frac{\mathrm{vol}_m(A)}{\mathrm{vol}_m(B)}$. Similarly, $N(A, B) \geq \frac{\gamma_m(A)}{\gamma_m(B)}$.*

Fact 43. *For any convex sets $A, B \subseteq \mathbb{R}^m$ with B symmetric, $N(A, B) \leq \frac{\mathrm{vol}_m(2A+B)}{\mathrm{vol}_m(B)}$.*

Fact 44. *For any convex sets $A, B \subseteq \mathbb{R}^m$ with A symmetric, $N(A, 2(A \cap B)) \leq N(A, B)$.*

Theorem 45 (Duality of Covering Numbers, [89])**.** *Given symmetric convex bodies $K, T \subseteq \mathbb{R}^n$, we have*

$$2^{-\Theta(n)} \cdot \mathcal{N}(T^\circ, K^\circ) \leq \mathcal{N}(K, T) \leq 2^{\Theta(n)} \cdot \mathcal{N}(T^\circ, K^\circ).$$

We show that a $2^{-O(n)}$ Gaussian measure lower bound is equivalent to a $2^{O(n)}$ upper bound for certain covering numbers.

Lemma 46. *The following conditions are equivalent for a symmetric convex body $D \subseteq \mathbb{R}^n$:*

1. $\gamma_n(D) \geq 2^{-O(n)}$,
2. $\mathcal{N}(\sqrt{n}B_2^n, D) \leq 2^{O(n)}$,
3. $\mathcal{N}(nB_1^n, D) \leq 2^{O(n)}$,
4. $\mathcal{N}(D^\circ, \frac{1}{\sqrt{n}}B_2^n) \leq 2^{O(n)}$,
5. $\mathcal{N}(D^\circ, \frac{1}{n}B_\infty^n) \leq 2^{O(n)}$.

Proof. We start by proving that condition (1) implies (2). Suppose $\gamma_n(D) \geq 2^{-O(n)}$, then Theorem 6 implies $\gamma_n(D') \geq 2^{-O(n)}$, where we define $D' := D \cap \sqrt{n}B_2^n$. We thus also have $\mathrm{vol}_n(D') \geq \gamma_n(D') \geq 2^{-O(n)}$. Then by Lemma 37, we have

$$\mathcal{N}(\sqrt{n}B_2^n, D) \leq \mathcal{N}(\sqrt{n}B_2^n, D') \leq 2^n \cdot \frac{\mathrm{vol}_n(\sqrt{n}B_2^n + D')}{\mathrm{vol}_n(D')} \leq 2^n \cdot \frac{\mathrm{vol}_n(2\sqrt{n}B_2^n)}{\mathrm{vol}_n(D')} \leq 2^{O(n)}.$$

We next show that condition (2) implies (1). Since $\gamma_n(\sqrt{n}B_2^n) = \Omega(1)$, we have $\gamma_n(x+D) \geq 2^{-O(n)}$ for some $x \in \mathbb{R}^n$. Lemma 4 then gives $\gamma_n(D) \geq \gamma_n(x+D) \geq 2^{-O(n)}$.

The implication $(3) \Rightarrow (2)$ immediately follows from $\sqrt{n}B_2^n \subseteq nB_1^n$. To prove the reverse implication $(2) \Rightarrow (3)$, we use Lemma 37 to obtain

$$\mathcal{N}(\sqrt{n}B_1^n, B_2^n) \leq 2^n \cdot \frac{\mathrm{vol}_n(\sqrt{n}B_1^n + B_2^n)}{\mathrm{vol}_n(B_2^n)} \leq 2^{O(n)} \cdot \frac{\mathrm{vol}_n(\sqrt{n}B_1^n)}{\mathrm{vol}_n(B_2^n)} \leq 2^{O(n)}.$$

It thus follows that $\mathcal{N}(nB_1^n, D) \leq \mathcal{N}(nB_1^n, \sqrt{n}B_2^n) \cdot \mathcal{N}(\sqrt{n}B_2^n, D) \leq 2^{O(n)}$.

The last two equivalences follow from the duality of covering numbers in Theorem 45.

□

As a corollary of Fact 43 we also have the following.

Lemma 47. *For symmetric convex sets $A, B \subseteq \mathbb{R}^m$, $N(A,B) \leq 3^m \cdot \frac{\mathrm{Vol}(A)}{\mathrm{Vol}(A \cap B)}$.*

Proof. Since $A \cap B$ is also symmetric, by Fact 43 we have

$$N(A,B) \leq N(A, A\cap B) \leq \frac{\mathrm{Vol}(2A + (A\cap B)))}{\mathrm{Vol}(A\cap B)} \leq \frac{\mathrm{Vol}(3A)}{\mathrm{Vol}(A\cap B)} = 3^m \cdot \frac{\mathrm{Vol}(A)}{\mathrm{Vol}(A\cap B)}. \quad \square$$

We will also need a more specific result about covering numbers of slices by subspaces.

Lemma 48. *For any symmetric convex $A, B \subseteq \mathbb{R}^m$ and any d-dimensional subspace $U \subseteq \mathbb{R}^m$, $N(A\cap U, B\cap U) \leq 6^d \cdot N(A,B)$.*

Proof. By Lemma 47, we have

$$N(A\cap U, B\cap U) \leq N(A\cap U, A\cap B\cap U) \leq 3^d \cdot \frac{\mathrm{Vol}(A\cap U)}{\mathrm{Vol}(A\cap B\cap U)}.$$

It remains to note that by Fact 44, we can cover A with $N(A,B)$ copies of $2(A\cap B)$, so that Fact 42 yields $\mathrm{Vol}(A\cap U) \leq N(A,B)\cdot \mathrm{Vol}(2(A\cap B)\cap U) = 2^d \cdot N(A,B)\cdot \mathrm{Vol}(A\cap B\cap U)$. $\square$

We are now ready to state and prove our main technical lemma.

Lemma 49. *Let $K \subseteq \mathbb{R}^m$ be a symmetric convex set such that $\gamma_m(\frac{C_0}{\alpha}K + \alpha\sqrt{m}B_2^m) \geq \frac{1}{2}$ for all $\alpha > 0$ and let $\mathcal{E}$ be a 1-regular M-ellipsoid for K. Then there is a universal constant C' such that for every $r > 0$, the number of axes of $\mathcal{E}$ of length at most r is at most $C'\sqrt{r}\cdot m^{\frac{3}{4}}$.*

Proof. Let U denote the span of directions corresponding to the axes of $\mathcal{E}$ of length at most r with $d := \dim(U)$. To simplify notation, let $K(\alpha) := \frac{C_0}{\alpha}K + \alpha\sqrt{m}B_2^m$, so that we still have $\gamma_U(K(\alpha)\cap U) \geq \gamma_m(K(\alpha)) \geq \frac{1}{2}$, where γ_U is the d-dimensional Gaussian measure in U. In particular, letting B_2^U denote the unit Euclidean ball restricted to U, $N(\sqrt{d}B_2^U, K(\alpha)\cap U) \leq \exp(C_2 d)$ by Fact 46.

Further, note that for any $\alpha > 0$ we have

$$\begin{aligned}
N\Big(K(\alpha)\cap U, \Big(\frac{C_0}{\alpha}\cdot\frac{mr}{d} + \alpha\sqrt{m}\Big)B_2^U\Big) &\leq N\Big(K(\alpha)\cap U, \Big(\frac{C_0}{\alpha}\cdot\frac{m}{d}\mathcal{E} + \alpha\sqrt{m}B_2^m\Big)\cap U\Big)\\
&\leq 6^d \cdot N\Big(K(\alpha), \frac{C_0}{\alpha}\cdot\frac{m}{d}\mathcal{E} + \alpha\sqrt{m}B_2^m\Big)\\
&\leq 6^d \cdot N(K, \frac{m}{d}\mathcal{E}) \leq 6^d \cdot \exp(C_1 d),
\end{aligned}$$

where in the second inequality we use Lemma 48, in the third we use Fact 41 and in the fourth we use Proposition 39. Setting $\alpha := \sqrt{C_0 \frac{mr}{d}} \cdot m^{-\frac{1}{4}}$, we have $\frac{C_0}{\alpha} \cdot \frac{mr}{d} + \alpha\sqrt{m} = \sqrt{C_0 \frac{r}{d}} \cdot m^{\frac{3}{4}}$, so by Fact 40,

$$N\left(\sqrt{d}B_2^U, 2\sqrt{\frac{C_0 r}{d}} m^{\frac{3}{4}} B_2^U\right) \le N\left(\sqrt{d}B_2^U, K(\alpha) \cap U\right) \cdot N\left(K(\alpha) \cap U, 2\sqrt{\frac{C_0 r}{d}} m^{\frac{3}{4}} B_2^U\right)$$
$$\le \exp(C_2 d) \cdot 6^d \cdot \exp(C_1 d).$$

On the other hand, by Fact 42,

$$N\left(\sqrt{d}B_2^U, 2\sqrt{\frac{C_0 r}{d}} m^{\frac{3}{4}} B_2^U\right) \ge \frac{\mathrm{Vol}_d(\sqrt{d}B_2^U)}{\mathrm{Vol}_d(2\sqrt{\frac{C_0 r}{d}} m^{\frac{3}{4}} B_2^U)} = \left(\frac{d}{2\sqrt{C_0 r} m^{\frac{3}{4}}}\right)^d.$$

Combining the above two displays yields the claim. □

Proof of Theorem 10. First we show a Gaussian measure lower bound for a 1-regular M-ellipsoid $\mathcal{E}$ of K with axes of lengths $\{\lambda_i\}_{i \in [m]}$ sorted in increasing order. Let k denote the maximum index with $\lambda_k \le \sqrt{m}$ and let $\mathcal{E}'$ denote the ellipsoid with the same eigenvectors as $\mathcal{E}$ and compressed axes of length $\min\{\lambda_i, \sqrt{m}\}$, so that in particular $\mathcal{E}' \subseteq \mathcal{E}$. Note that

$$\begin{aligned}
\gamma_m(\mathcal{E}) &\ge \gamma_m(\mathcal{E}') \\
&\ge \int_{\mathcal{E}'} \frac{1}{(2\pi)^{m/2}} \exp(-\tfrac{1}{2} \underbrace{\|x\|_2^2}_{\le m}) \ \mathrm{d}x \\
&\ge \exp(-C'm) \cdot \mathrm{Vol}(\mathcal{E}') \\
&\ge \exp(-C''m) \prod_{i \in [k]} \frac{\lambda_i}{\sqrt{m}},
\end{aligned}$$

for some constants $C', C'' > 0$. Denote $I_r := \{i \in [k] : \lambda_i \in [\frac{r}{2}, r]\}$. We apply Lemma 49 as follows:

$$\prod_{i \in [k]} \frac{\lambda_i}{\sqrt{m}} \ge \prod_{r \in \sqrt{m} \cdot 2^{\mathbb{Z}_{\le 0}}} \left(\frac{r}{2\sqrt{m}}\right)^{|I_r|} \ge \prod_{r \in \sqrt{m} \cdot 2^{\mathbb{Z}_{\le 0}}} \left(\frac{r}{2\sqrt{m}}\right)^{C'\sqrt{r} m^{\frac{3}{4}}}.$$

Denoting $\frac{r}{2\sqrt{m}} = 2^{-a}$, the product becomes, for a constant C_3,

$$\prod_{i \in [k]} \frac{\lambda_i}{\sqrt{m}} \ge \left(\prod_{a=0}^{\infty} (2^{-a})^{2^{-\frac{a}{2}}}\right)^{C'm} = \exp(-C_3 m).$$

Finally, we use Fact 42 and Proposition 39 to lower bound the Gaussian measure of K:

$$\gamma_m(K) \geq \frac{\gamma_m(\mathcal{E})}{N(\mathcal{E}, K)} \geq \exp(-(C'' + C_1 + C_3)m). \qquad \square$$

1.6 One-dimensional Gaussian measure bounds

Proof of Lemma 15. We make use of the following tail inequality due to Szarek and Werner [161] which holds for $t > -1$:

$$\Pr_{g \sim N(0,1)}[g > t] < \frac{1}{\sqrt{2\pi}} \frac{4e^{-t^2/2}}{3t + (t^2 + 8)^{1/2}}.$$

In particular, for $t \geq 1$ the right side is upper bounded by $\frac{1}{\sqrt{2\pi}} \frac{4e^{-t^2/2}}{6}$. Thus

$$\Pr_{g \sim N(0,1)}[|g| \leq t] \geq 1 - \frac{4}{3\sqrt{2\pi}} e^{-t^2/2}.$$

Since the function $z \mapsto e^{-2z/3}$ is convex, we have $1 - \frac{4}{3\sqrt{2\pi}} z \geq e^{-2z/3}$ for all $z \in [0, e^{-1/2}]$ as it holds for the endpoints of the interval. Therefore for $t \geq 1$,

$$\Pr_{g \sim N(0,1)}[|g| \leq t] \geq \exp(-\tfrac{2}{3} e^{-t^2/2}).$$

We conclude that for any $a \in \mathbb{R}^n$ with $\|a\|_2 \leq 1$ and $t \geq 1$ one has

$$\Pr_{y \sim N(0,I_n)}[|\langle a, y \rangle| \leq t] = \Pr_{g \sim N(0,1)}\left[|g| \leq \frac{t}{\|a\|_2}\right] \geq \exp(-\tfrac{2}{3} e^{-t^2/(2\|a\|_2^2)}) \geq \exp(-e^{-t^2/2} \cdot \|a\|_2^2).$$

Indeed, the last inequality follows because

$$\frac{2}{3} \exp\left(\frac{t^2}{2} - \frac{t^2}{2\|a\|_2^2}\right) \leq \frac{2}{3} \exp\left(\frac{1}{2} - \frac{1}{2\|a\|_2^2}\right) \leq \frac{2}{3} \cdot e^{1/2} \cdot \frac{2}{e} \cdot \|a\|_2^2 \leq \|a\|_2^2,$$

where the second to last inequality follows from $e^z \geq ez$ for $z := 1/(2\|a\|_2^2)$. $\square$

Proof of Lemma 21. Draw another random variable $z \sim N(0, B - A)$ and note that by log-concavity we have

$$\Pr_{y \sim N(0,A)}[y \in K] \geq \Pr_{y \sim N(0,A)}\left[\Pr_{z \sim N(0,B-A)}[y + z \in K]\right] = \Pr_{y \sim N(0,B)}[y \in K]. \qquad \square$$

1.7 Large convex sets without partial colorings

We have mentioned earlier that a symmetric convex set K with measure $\gamma_n(K) \geq e^{-\delta n}$ contains a partial coloring $\boldsymbol{x} \in \{-1,0,1\}^n$ with a linear number of nonzero coordinates if the constant δ is small enough — but this is false for constants beyond a certain threshold, even if one is allowed to rescale the body by some parameter dependent on δ. The construction for such a set is a thin strip that avoids any point in $\{-1,0,1\}^n \setminus \{\mathbf{0}\}$.

Lemma 50. *For any $C \geq 1$, there exists a $\delta > 0$ so that the following holds: for any $n \in \mathbb{N}$ large enough there is a symmetric convex body $K \subseteq \mathbb{R}^n$ so that (i) $(C^n K) \cap (\{-1,0,1\}^n \setminus \{\mathbf{0}\}) = \emptyset$ and (ii) $\gamma_n(K) \geq e^{-\delta n}$.*

Proof. The construction is probabilistic. We sample a Gaussian $\boldsymbol{g} \sim N(\mathbf{0}, \boldsymbol{I}_n)$ and for a tiny parameter $s > 0$ that we determine later, we consider the strip $K := \{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{g}, \boldsymbol{x}\rangle| \leq s\}$. Consider the set of nontrivial partial colorings $X := \{-1,0,1\}^n \setminus \{\mathbf{0}\}$ and recall that $|X| \leq 3^n$. For any $\boldsymbol{x} \in X$, the distribution of $\langle \boldsymbol{g}, \boldsymbol{x}\rangle$ is Gaussian with variance $\|\boldsymbol{x}\|_2^2 \geq 1$ and hence the density of this 1-dimensional Gaussian is at most $\frac{1}{\sqrt{2\pi}} e^0 \leq \frac{1}{2}$ everywhere. In particular for a fixed $\boldsymbol{x} \in X$, one can obtain the simple estimate of $\Pr[|\langle \boldsymbol{g}, \boldsymbol{x}\rangle| \leq t] \leq 4t$ for any $t > 0$. Then choosing $s := \frac{1}{16} \cdot C^{-n} 3^{-n}$ we obtain

$$\Pr_{\boldsymbol{g}}\left[(C^n K) \cap X \neq \emptyset\right] \leq \sum_{\boldsymbol{x} \in X} \Pr_{\boldsymbol{g}}\left[|\langle \boldsymbol{g}, \boldsymbol{x}\rangle| > C^n s\right] \leq \frac{1}{4} \cdot |X| \cdot 3^{-n} \leq \frac{1}{4} \qquad (*)$$

Moreover using Markov's Inequality we obtain the (rather weak) estimate

$$\Pr\left[\|\boldsymbol{g}\|_2^2 > 4n\right] \leq \frac{1}{4} \qquad (**)$$

Then with probability at least $1/2$ none of the events $(*)$ and $(**)$ happen. We fix such an outcome of $\boldsymbol{g}$ and estimate that the measure of our strip is

$$\gamma_n(K) = \int_{-s/\|\boldsymbol{g}\|_2}^{s/\|\boldsymbol{g}\|_2} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \,\mathrm{d}x \geq \frac{1}{\sqrt{2\pi}} e^{-1/2} \frac{2s}{\sqrt{n}} \geq e^{-\delta n}$$

for a suitable choice of δ using $\frac{s}{\|\boldsymbol{g}\|_2} \leq 1$. $\square$

Chapter 2

VECTOR BALANCING IN LEBESGUE SPACES

This chapter is based on a joint paper with Thomas Rothvoss [140].

2.1 Introduction

The celebrated *Spencer's Theorem* in discrepancy theory [157] shows that "six standard deviations suffice" for balancing vectors in the ℓ_∞-norm: for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in [-1,1]^n$, there exist signs $\boldsymbol{x} \in \{-1,1\}^n$ such that $\|\sum_{i=1}^n x_i \boldsymbol{a}_i\|_\infty \leq 6\sqrt{n}$. More generally, Spencer showed that for vectors in $[-1,1]^d$ with $n \leq d$ one can achieve a bound of $O(\sqrt{n \log(2d/n)})$. While his proof used a nonconstructive form of the *partial coloring lemma* based on the pigeonhole principle, in the past decade several approaches starting with the breakthrough work of Bansal [16] did succeed in computing such signs in polynomial time [103, 146, 100, 61].

As for balancing vectors of bounded ℓ_2-norm, the situation has been more delicate. In the same paper, Spencer [157] showed a nonconstructive bound of $O(\log n)$ for the ℓ_∞ discrepancy of vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_2^d$ and also stated a discrete version of a conjecture of Komlós that this may be improved to $O(1)$. This was improved to $O(\sqrt{\log n})$ by Banaszczyk [13] who showed that in fact for any set of n vectors of ℓ_2-norm at most 1 and any convex body $K \subseteq \mathbb{R}^d$ of Gaussian measure at least $1/2$, some ± 1 combination of such vectors lies in $5 \cdot K$. For the general setting of ℓ_q discrepancy, Matoušek [110] gave an upper bound of $O(q) \cdot d^{1/q}$ for balancing vectors from ℓ_2 to ℓ_q. More recently, the work of Barthe, Guédon, Mendelson and Naor [30] (see Prop. 25) shows that, for $q \geq 2$, n-dimensional slices of the ℓ_q ball in $\mathbb{R}^d$ scaled by a factor of $O(\sqrt{q}) \cdot n^{1/q}$ do have Gaussian measure at least $1/2$, thus improving the bound to $O(\sqrt{q}) \cdot n^{1/q}$. For $q = \log n$, this matches the ℓ_2 to ℓ_∞ bound of $O(\sqrt{\log n})$. Banaszczyk's proof was nonconstructive and the first polynomial time algorithm in the general convex body setting was found only recently by Bansal, Dadush, Garg and Lovett [19], while the Komlós conjecture remains an open problem. The work of [19] actually shows that for any vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_2^d$ there exists an efficiently computable distribution over signs $\boldsymbol{x} \in \{-1, 1\}^n$ so that the sum $\boldsymbol{X} := \sum_{i=1}^n x_i \boldsymbol{a}_i$ is $O(1)$-subgaussian, meaning that $\mathbb{E}[e^{\langle \boldsymbol{\theta}, \boldsymbol{X} \rangle}] \leq e^{O(1)\|\boldsymbol{\theta}\|_2^2}$ for every $\boldsymbol{\theta} \in \mathbb{R}^d$, and will be in $O(1) \cdot K$ with good probability. Interestingly, this means their algorithm is *oblivious* to the body K, which is a striking difference to the regime of $\gamma_n(K) = e^{-\Theta(n)}$ where any algorithm needs to be dependent on K. The connection between Banaszczyk's theorem and subgaussianity is due to Dadush et al. [51].

For the general setting of balancing vectors from ℓ_p to ℓ_q, where we are given vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$ and wish to find signs $x_1, \ldots, x_n$ that minimize the ℓ_q norm of $\sum_{i=1}^n x_i \boldsymbol{a}_i$ (also called ℓ_q discrepancy), not much was known beyond Spencer's theorem ($p = \infty$) or what can be deduced from Banaszczyk's theorem as above: any vector in B_p^d also belongs to $d^{\max(0,1/2-1/p)} \cdot B_2^d$, thus implying a discrepancy bound of $O(\sqrt{q}) \cdot d^{\max(0,1/2-1/p)} \cdot n^{1/q}$. Even in the square case $d = n$, in spite of tight partial coloring bounds [157], it has been an open problem to remove the dependency on $\sqrt{q}$ [53]. The goal of this paper is to provide a unified approach for balancing from ℓ_p to ℓ_q via optimal constructive fractional partial colorings, which yield optimal bounds for most of the range $1 \leq p \leq q \leq \infty$. We obtain such fractional partial colorings by proving a new measure lower bound on the relevant linear preimages of ℓ_q balls (Section 3) and an improved algorithm for sets of Gaussian measure $e^{-\delta n}$ for any $\delta > 0$ (Section 4), as opposed to previous work ([146, 61]) which required measure $e^{-\delta n}$ for *sufficiently small* $\delta > 0$. Finally, we show that a *hereditary* volume lower bound is sufficient to imply such Gaussian measure bound (Section 5).

As an application, we show a slight improvement to the bounds for the well-known Beck-Fiala conjecture [33], a discrete version of Komlós. It asks for a $O(\sqrt{t})$ bound on the ℓ_∞ discrepancy of any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \{0, 1\}^d$, each with at most t ones. We establish the conjecture for $t \geq n$ and show slightly improved bounds when t is close to n (Corollary 54).

Notation. Let $B_p^d := \{\boldsymbol{x} \in \mathbb{R}^d : \|\boldsymbol{x}\|_p \leq 1\}$ denote the unit ball in the ℓ_p-norm. The *Gaussian measure* of a measurable set $K \subseteq \mathbb{R}^n$ is given by $\gamma_n(K) := \Pr_{\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)}[\boldsymbol{x} \in K]$. We denote the *mean width* of a convex set as $w(K) := \mathbb{E}_{\boldsymbol{\theta} \in S^{n-1}}[\sup_{\boldsymbol{x} \in K} \langle \boldsymbol{\theta}, \boldsymbol{x} \rangle]$. The Euclidean distance to a set $S \subseteq \mathbb{R}^n$ is denoted by $d(\boldsymbol{x}, S) := \min\{\|\boldsymbol{x} - \boldsymbol{y}\|_2 : \boldsymbol{y} \in S\}$. A function $f : \mathbb{R}^d \to \mathbb{R}$ is α-Lipschitz if $|f(\boldsymbol{x}) - f(\boldsymbol{y})| \leq \alpha \cdot \|\boldsymbol{x} - \boldsymbol{y}\|_2$ for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d$. If $\boldsymbol{A} \in \mathbb{R}^{d \times n}$ is a matrix, we denote its rows by $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_d \in \mathbb{R}^n$ and its columns by $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \mathbb{R}^d$. Naturally, a matrix can also be interpreted as a (not necessarily invertible) linear map. Then for any set $K \subseteq \mathbb{R}^d$, we use the notation $\boldsymbol{A}^{-1}(K) := \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{A}\boldsymbol{x} \in K\}$. The C-scaling of a symmetric convex body K is the body $C \cdot K = \{c\boldsymbol{x} : \boldsymbol{x} \in K\}$.

2.1.1 *Our contribution*

Our main contribution is a tight bound on partial colorings for balancing from ℓ_p to ℓ_q:

Theorem 51. *Let $n \leq d$ and $2 \leq p \leq q \leq \infty$.*[1] *Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$, there exists a polynomial-time computable partial coloring $\boldsymbol{x} \in [-1,1]^n$ with $|\{i : x_i^2 = 1\}| \geq n/2$ so that*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \lesssim \sqrt{\min\left(p, \log\left(\frac{2d}{n}\right)\right)} \cdot n^{1/2 - 1/p + 1/q}.$$

By Theorem 3, the condition $n \leq d$ does not weaken the theorem: in fact for $n > d$ the upper bound can only be larger than that of $n = d$ by a factor of two. On the other hand, the condition $p \leq q$ is natural, for otherwise if $p > q$ we would need a polynomial dependence on the dimension d, even for $n = 1$. By iteratively applying Theorem 51 we can obtain a full coloring at the expense of another factor of $\frac{1}{1/2 - 1/p + 1/q}$, with the caveat that $p > 2$ whenever $q = \infty$:

Theorem 52. *Let $n \leq d$ and $2 \leq p \leq q \leq \infty$ with $\{p, q\} \neq \{2, \infty\}$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in$*

[1]When $p \leq 2$, uniformly random signs achieve a tight bound of $\Theta(n^{1/q})$ (see Theorem 55), so we focus on the more interesting case $p \geq 2$.

B_p^d, *there exist polynomial-time computable signs* $\boldsymbol{x} \in \{-1,1\}^n$ *so that*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \lesssim \frac{\sqrt{\min\left(p, \log\left(\frac{2d}{n}\right)\right)}}{1/2 - 1/p + 1/q} \cdot n^{1/2-1/p+1/q}.$$

This significantly improves upon the general $\sqrt{q} \cdot d^{1/2-1/p} \cdot n^{1/q}$ bound from Banaszczyk's theorem in [53] when $p = 2 + \varepsilon$ for (not too small) $\varepsilon > 0$ and $q \gg 1$. It is also worth noting that we may always assume $q \leq \log d$ as larger norms are equivalent up to a constant by Lemma 56. When $p = q$ and $d = n$, we get the following corollary which matches, up to a constant, the lower bound $\Omega(\sqrt{n})$ of [11] known to hold for any norm:

Corollary 53 (ℓ_p version of Spencer's theorem)**.** *Let* $2 \leq p \leq \infty$ *and* $n \in \mathbb{N}$. *Then for any* $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^n$, *there exist polynomial-time computable signs* $\boldsymbol{x} \in \{-1,1\}^n$ *so that*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_p \lesssim \sqrt{n}.$$

The following corollary shows the Beck-Fiala conjecture holds for $t \geq n$ and slightly improves upon the best known bound of $O(\sqrt{t \log n})$ [13] when t is close to n:

Corollary 54 (Bound for Beck-Fiala)**.** *Let* $n \leq d$ *and* $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \{0,1\}^d$, *each with at most* $t \in [d]$ *ones. Then there exist polynomial-time computable signs* $\boldsymbol{x} \in \{-1,1\}^n$ *so that*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_\infty \lesssim \sqrt{t} \log\left(\frac{2\max(n,t)}{t}\right).$$

We show the partial coloring bound in Theorem 51 is tight at least when $d = n$:

Theorem 55. *Let* $1 \leq p \leq q \leq \infty$. *There exist infinitely many positive integers* n *for which we can find* $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^n$ *such that for any* $\boldsymbol{x} \in [-1,1]^n$ *with* $|\{i : x_i^2 = 1\}| \geq n/2$ *one has*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \gtrsim n^{\max(0, 1/2-1/p)+1/q}.$$

2.2 *Preliminaries*

We will use two elementary inequalities dealing with ℓ_p-norms. The first one estimates the ratio between different norms:

Lemma 56. *For any $\boldsymbol{z} \in \mathbb{R}^d$ and $1 \leq p \leq q \leq \infty$, we have $\|\boldsymbol{z}\|_q \leq \|\boldsymbol{z}\|_p \leq m^{1/p-1/q}\|\boldsymbol{z}\|_q$.*

It is instructive to note that this bound implies $\|\boldsymbol{z}\|_\infty \leq \|\boldsymbol{z}\|_{\log_2(d)} \leq 2\|\boldsymbol{z}\|_\infty$. If one has an upper bound on the largest entry in a vector — say $\|\boldsymbol{z}\|_\infty \leq 1$ — then one can strengthen the first inequality to $\|\boldsymbol{z}\|_q^q \leq \|\boldsymbol{z}\|_p^p$. More generally:

Lemma 57. *For any $\boldsymbol{z} \in \mathbb{R}^d$ and $1 \leq p \leq q \leq \infty$, we have $\|\boldsymbol{z}\|_q^q \leq \|\boldsymbol{z}\|_p^p \cdot \|\boldsymbol{z}\|_\infty^{q-p}$.*

We will also need the following version of *Khintchine's inequality*, see e.g. the excellent textbook of Artstein-Avidan, Giannopoulos and Milman [7].

Lemma 58 (Khintchine's inequality)**.** *Given $p > 0$, $a_1, \dots, a_n \in \mathbb{R}$ and $\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)$, we have*

$$\mathbb{E}\Big[\Big|\sum_{i=1}^n x_i a_i\Big|^p\Big] \lesssim \sqrt{p} \cdot \Big(\sum_{i=1}^n a_i^2\Big)^{p/2}.$$

This fact can be derived from a standard Chernov bound which guarantees that for a vector with $\|\boldsymbol{a}\|_2 = 1$ one has $\Pr[|\langle \boldsymbol{a}, \boldsymbol{x}\rangle| > \lambda] \leq 2e^{-\lambda^2/2}$; then one can analyze that the regime of $\lambda = \Theta(\sqrt{p})$ dominates the contribution to $\mathbb{E}[|\langle \boldsymbol{a}, \boldsymbol{x}\rangle|^p]$. We use it to show the following standard estimate on the type constants of ℓ_p spaces (see Section 2.7):

Lemma 59. *Given $p \geq 1$ and $\boldsymbol{a}_1, \dots, \boldsymbol{a}_n \in B_p^d$ and $\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)$, we have*

$$\mathbb{E}\Big[\Big\|\sum_{i=1}^n x_i \boldsymbol{a}_i\Big\|_p\Big] \lesssim \sqrt{p} \cdot n^{\max(1/2, 1/p)}.$$

2.3 *Main technical result*

In this section we show our measure lower bound for balancing vectors from ℓ_p to ℓ_q:

Theorem 60. *Let $n \le d$ and $1 \le p \le q \le \infty$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$,*

$$\gamma_n\Big(\Big\{\boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \le \sqrt{\min\Big(p, \log\Big(\frac{2d}{n}\Big)\Big)} \cdot n^{\max(0,1/2-1/p)+1/q}\Big\}\Big) \ge 2^{-O(n)}.$$

In order to show Theorem 60, roughly speaking it will suffice to show the corresponding bounds for the two special cases of $q \in \{p, \infty\}$, which can be bootstrapped into a general bound. First we address the simpler case $p = q$ which at heart is based on Khintchine's inequality:

Lemma 61. *Let $n \le d$ and $p \ge 1$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$,*

$$\gamma_n\Big(\Big\{\boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_p \le \sqrt{p} \cdot n^{\max(1/2,1/p)}\Big\}\Big) \ge 2^{-O(n)}.$$

Proof. By Lemma 59 we know that, for some constant $C > 0$,

$$\mathop{\mathbb{E}}_{\boldsymbol{x} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)} \Big[\Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_p\Big] \le C\sqrt{p} \cdot n^{\max(1/2,1/p)}.$$

By Markov's inequality it follows that

$$\gamma_n\Big(\Big\{\boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_p \le 2C\sqrt{p} \cdot n^{\max(1/2,1/p)}\Big\}\Big) \ge 1/2,$$

so that the result follows by Lemma 16. $\square$

Next, we deal with the crucial case $q = \infty$:

Lemma 62. *Let $n \le d$ and $p \ge 1$. Then for any $\boldsymbol{A} \in \mathbb{R}^{d \times n}$ with columns $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$ and rows $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_d \in \mathbb{R}^n$, the body $K := \{\boldsymbol{x} \in \mathbb{R}^n : \| \sum_{i=1}^n x_i \boldsymbol{a}_i \|_\infty \le \sqrt{p} \cdot n^{\max(0,1/2-1/p)}\}$ satisfies*

$$\gamma_n(K) \ge \prod_{j \in [d]} \gamma_n(\{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{x}, \boldsymbol{A}_j \rangle| \le \sqrt{p} n^{\max(0,1/2-1/p)}\}) \ge 2^{-O(n)}.$$

Proof. The main idea in the proof is that we can convert the bound on the ℓ_p-norm of the columns $\boldsymbol{a}_i$ into information about the ℓ_2-norm of the rows $\boldsymbol{A}_j$. Namely,

$$\Big(\frac{1}{n} \sum_{j \in [d]} \|\boldsymbol{A}_j\|_2^p\Big)^{1/p} \overset{\text{Lem 56}}{\le} n^{\max(0,1/2-1/p)} \cdot \Big(\frac{1}{n} \underbrace{\sum_{j \in [d]} \|\boldsymbol{A}_j\|_p^p}_{\le n}\Big)^{1/p} \le n^{\max(0,1/2-1/p)}. \tag{2.1}$$

We rescale the row vectors to $\boldsymbol{V}_j := (\sqrt{p}n^{\max(0,1/2-1/p)})^{-1}\boldsymbol{A}_j$ and abbreviate $y_j := \|\boldsymbol{V}_j\|_2^2$, so that Eq. (2.1) simplifies to $\sum_{j=1}^d y_j^{p/2} \leq n \cdot p^{-p/2}$. We may then apply Šidak's Lemma 5 and bound the one-dimensional measure:

$$
\begin{aligned}
\gamma_n(K) &= \gamma_n\big(\big\{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{x}, \boldsymbol{V}_j\rangle| \leq 1 \;\; \forall j \in [d]\big\}\big) \\
&\overset{\text{Lem 5}}{\geq} \prod_{j\in[d]} \gamma_n\big(\big\{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{x}, \boldsymbol{V}_j\rangle| \leq 1\big\}\big) \\
&\overset{\text{Lem 14}}{\geq} \prod_{j\in[d]} \big(1 - \exp(-y_j^{-1}/2)\big) \\
&\overset{\text{Claim I}}{\geq} \prod_{j\in[d]} \exp\Big(-C'p^{p/2}y_j^{p/2}\Big) = \exp\Big(-C'p^{p/2}\sum_{j\in[d]} y_j^{p/2}\Big) \geq \exp(-C'n)
\end{aligned}
$$

Here we have used an estimate that remains to be proven:

Claim I. *For any $p \geq 1$ and $y > 0$ one has $1 - \exp(-\frac{1}{2y}) \geq \exp(-C'p^{p/2}y^{p/2})$ where $C' > 0$ is a universal constant.*

Proof of Claim I. It will suffice to show for any $y > 0$:

$$-\log(1 - \exp(-y^{-1}/2)) \leq O(p^{p/2}y^{p/2}).$$

To see this, let $z = \sqrt{2y}$ and note that it suffices to show

$$-\log(1 - \exp(-z^{-2})) \cdot z^{-p} \leq O((p/2)^{p/2}).$$

First, by convexity of $x \mapsto -\log(1-x)$, we have $-\log(1-x) \leq O(x)$ for $x \in [0, 1/e]$. It follows that for $z \leq 1$, we have

$$-\log(1 - \exp(-z^{-2})) \leq O(\exp(-z^{-2})) \leq O(\lceil p/2\rceil!/z^{-2\lceil p/2\rceil}),$$

and therefore $-\log(1 - \exp(-z^{-2})) \cdot z^{-p} \leq O(\lceil p/2\rceil!) \leq O((p/2)^{p/2})$.

Next, we claim that $-\log(1 - \exp(-z^{-2})) \leq 4z$ for all $z > 0$. Indeed, both sides tend to 0 as $z \to 0$ and the derivative of the left side is

$$\frac{2}{z^3\Big(\exp\Big(\frac{1}{2z^2}\Big) - 1\Big)} < \frac{2}{z^3\Big(\frac{1}{2z^2} + \frac{1}{8z^4}\Big)} = \frac{16z}{4z^2+1} \leq 4,$$

where we used $e^x > 1 + x + x^2/2$ for $x = \frac{1}{2z^2}$ and $(2z-1)^2 \geq 0$. It follows that when $z \geq 1$, $-\log(1 - \exp(-z^{-2})) \cdot z^{-p} \leq 4z^{1-p} \leq 4 \leq O((p/2)^{p/2})$. $\square$

Remark 3. *This argument is largely motivated by the result of Ball and Pajor [10] which bounds volume instead of Gaussian measure. More specifically, [10] prove that for $1 \leq p \leq \infty$ and any matrix $\boldsymbol{A} \in \mathbb{R}^{d \times n}$, the set*

$$K = \Big\{ \boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{A}_j, \boldsymbol{x} \rangle| \leq \sqrt{p} \cdot \Big(\frac{1}{n} \sum_{j=1}^{d} \|\boldsymbol{A}_j\|_2^p \Big)^{1/p} \ \forall j \in [d] \Big\}$$

satisfies $\mathrm{vol}_n(K) \geq 1$. In contrast, our Lemma 62 provides a simpler proof of a stronger result (up to a constant scaling), since the volume of a convex body is always at least its Gaussian measure. On the other hand, it is also possible to recover Lemma 62 directly from this result together with Theorem 9.

We are now ready to show Theorem 60:

Proof of Theorem 60. Let $1 \leq p \leq q \leq \infty$ and let $\boldsymbol{A} \in \mathbb{R}^{d \times n}$ denote the matrix with columns $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$. By Lemma 57 we know that for any $\boldsymbol{z} \in \mathbb{R}^d$ with $\|\boldsymbol{z}\|_p \leq n^{1/p}$ and $\|\boldsymbol{z}\|_\infty \leq 1$ one has $\|\boldsymbol{z}\|_q \leq (\|\boldsymbol{z}\|_p^p \cdot \|\boldsymbol{z}\|_\infty^{q-p})^{1/q} \leq n^{1/q}$. Phrased in geometric terms this means $n^{1/q} B_q^d \supseteq n^{1/p} B_p^d \cap B_\infty^d$. We would like to point out that this is a crucial point to obtain a dependence solely on n rather than the larger parameter d. Next, note the fact that $\boldsymbol{A}^{-1}(S \cap T) = \boldsymbol{A}^{-1}(S) \cap \boldsymbol{A}^{-1}(T)$ for any sets S and T which we use together with the inequality of Šidak and Kathri (Lemma 5) to obtain the estimate

$$\begin{aligned}
&\gamma_n\Big(\boldsymbol{A}^{-1}\big(\sqrt{p} \cdot n^{\max(0,1/2-1/p)+1/q} B_q^d\big)\Big) \\
&\geq \gamma_n\Big(\boldsymbol{A}^{-1}\big(\sqrt{p} \cdot n^{\max(0,1/2-1/p)}(n^{1/p} B_p^d \cap B_\infty^d)\big)\Big) \\
&\geq \gamma_n\Big(\boldsymbol{A}^{-1}\big(\sqrt{p} \cdot n^{\max(1/2,1/p)} B_p^d\big)\Big) \cdot \prod_{j \in [d]} \gamma_n\big(\{\boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{x}, \boldsymbol{A}_j \rangle| \leq \sqrt{p} n^{\max(0,1/2-1/p)}\}\big) \\
&\geq 2^{-O(n)} \cdot 2^{-O(n)} = 2^{-O(n)},
\end{aligned}$$

where we have used the measure lower bounds from Lemmas 61 and 62. This shows the claimed bound whenever $p \leq O(\log(\frac{2d}{n}))$, where the hidden constant can be removed by scaling the corresponding convex body, see Lemma 16.

It remains to prove that we can bootstrap the existing bound for the regime of large p. So let us assume that $p \geq 2 \cdot \max\{1, \log(d/n)\}$. Let $p_0 \in [2, p]$ be a parameter to be determined and remark that Lemma 56 gives $\|\boldsymbol{a}_i\|_{p_0} \leq d^{1/p_0 - 1/p} \cdot \|\boldsymbol{a}_i\|_p \leq d^{1/p_0-1/p}$. Applying the above measure lower bound for p_0 implies

$$\gamma_n\Big(\Big\{\boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \leq \sqrt{p_0} \cdot n^{1/2-1/p_0+1/q} \cdot d^{1/p_0-1/p}\Big\}\Big) \geq 2^{-O(n)}.$$

We can rewrite the above upper bound on ℓ_q-norm as

$$\sqrt{p_0} \cdot n^{1/2-1/p_0+1/q} \cdot d^{1/p_0-1/p} = n^{1/2-1/p+1/q} \cdot \underbrace{\Big(\frac{d}{n}\Big)^{-1/p}}_{\leq 1} \cdot \sqrt{p_0} \cdot \Big(\frac{d}{n}\Big)^{1/p_0}.$$

Taking $p_0 := 2 \cdot \max\{1, \log(d/n)\}$ gives the desired result as then $(d/n)^{1/p_0} \leq \sqrt{e}$ and Lemma 16 can again deal with such constant scaling. $\square$

Now our main result on existence of partial colorings easily follows:

Proof of Theorem 51. Apply Theorem 7 to the set

$$K := \Big\{\boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{a}_i \Big\|_q \leq \sqrt{\min\Big(p, \log\Big(\frac{2d}{n}\Big)\Big)} \cdot n^{1/2-1/p+1/q}\Big\},$$

which by Theorem 60 indeed has a Gaussian measure of $\gamma_n(K) \geq 2^{-O(n)}$. $\square$

Next, we show how to obtain a full coloring by iteratively finding partial colorings.

Proof of Theorem 52. Let again $2 \leq p \leq q \leq \infty$ and let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$. We begin with $\boldsymbol{x}^{(0)} := \boldsymbol{0}$ and given $\boldsymbol{x}^{(0)}, \ldots, \boldsymbol{x}^{(t)}$ we set $S^{(t)} := \{i \in [n] : -1 < x_i^{(t)} < 1\}$ as the *active variables*. Then combining Theorem 7 and Theorem 60 we can find a partial coloring $\boldsymbol{x}^{(t+1)} \in [-1, 1]^n$ in polynomial time so that $|S^{(t+1)}| \leq |S^{(t)}|/2$ and $\| \sum_{i=1}^n (x_i^{(t+1)} - x_i^{(t)}) \boldsymbol{a}_i \|_q \leq C_1 \sqrt{\min(p, \log(\frac{2d}{|S^{(t)}|}))} \cdot |S^{(t)}|^{1/2-1/p+1/q}$. Let $\boldsymbol{x}^{(T)}$ be the first iterate with $\boldsymbol{x}^{(T)} \in \{-1, 1\}^n$.

Clearly $|S^{(t)}| \leq n2^{-t}$ and $T \leq \log_2(n)$. Using the triangle inequality we get

$$\begin{aligned}\Big\| \sum_{i=1}^{n} x_i^{(T)} \boldsymbol{a}_i \Big\|_q &\leq \sum_{t=0}^{T-1} \Big\| \sum_{i=1}^{n} (x_i^{(t+1)} - x_i^{(t)}) \boldsymbol{a}_i \Big\|_q \\ &\leq C_1 \sum_{t=0}^{T-1} \sqrt{\min\Big(p, \log\Big(\frac{2d}{2^{-t} \cdot n}\Big)\Big)} \cdot (2^{-t} \cdot n)^{1/2-1/p+1/q} \\ &\leq \frac{C_1 C_2 \sqrt{\min\Big(p, \log\Big(\frac{2d}{n}\Big)\Big)}}{1/2 - 1/p + 1/q} \cdot n^{1/2-1/p+1/q}. \qquad \square\end{aligned}$$

The intuition behind the extra factor for obtaining a full coloring is as follows: abbreviate the exponent as $\beta := 1/2 - 1/p + 1/q$. Then it takes $\frac{1}{\beta}$ iterations until the term $|S^{(t)}|^\beta$ decreases by a factor of 1/2 which dominates the miniscule growth of the logarithmic term. Then indeed the overall discrepancy is dominated by the discrepancy from the first $\frac{1}{\beta}$ iterations.

We can now demonstrate how a nontrivial choice of ℓ_p-norms can be beneficial in classical discrepancy settings:

Proof of Corollary 54. Consider columns $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \{0,1\}^d$ with at most t nonzero entries per $\boldsymbol{a}_i$. First let us study the case $t \geq n/10$. Since for each column $\|\boldsymbol{a}_i\|_4 \leq t^{1/4}$, Theorem 52 provides a coloring $\boldsymbol{x} \in \{-1,1\}^n$ with $\| \sum_{i=1}^n x_i \boldsymbol{a}_i \|_\infty \leq O(n^{1/4} t^{1/4}) = O(\sqrt{t})$. [2]

Now if $t < n/10$, we take $p \in [2, 16)$ with $1/2 - 1/p = 1/\log(n/t)$. Then $\|\boldsymbol{a}_i\|_p \leq t^{1/p}$ and Theorem 52 gives $\boldsymbol{x} \in \{-1,1\}^n$ with

$$\Big\| \sum_{i=1}^{n} x_i \boldsymbol{a}_i \Big\|_\infty \leq \frac{C \cdot n^{1/2-1/p} \cdot t^{1/p}}{1/2 - 1/p} = C\sqrt{t} \log(n/t) \cdot \underbrace{(n/t)^{1/\log(n/t)}}_{=e}. \qquad \square$$

We conclude this section by showing that the term $n^{\max(0,1/2-1/p)+1/q}$ in our bounds is necessary:

Proof of Theorem 55. Consider the case $p \geq 2$. Consider an $n \times n$ *Hadamard matrix*, which is a matrix $\boldsymbol{H} \in \{-1,1\}^{n \times n}$ so that all rows and columns are orthogonal. Such matrices

[2] In fact for $t \geq n$ a more careful choice of $p = \log(2t/n)$ gives a better ℓ_∞ discrepancy bound of $O(\sqrt{n \log(2t/n)})$, even though the Beck-Fiala conjecture asks only for $O(\sqrt{t})$.

are known to exist at least whenever n is a power of 2. The columns satisfy $\|\boldsymbol{h}_i\|_p = n^{1/p}$ and for any $\boldsymbol{x} \in [-1,1]^n$ with $|\{i : x_i^2 = 1\}| \geq n/2$ we know that $\|\boldsymbol{x}\|_2 \geq \Omega(\sqrt{n})$ and $\|\boldsymbol{H}\boldsymbol{x}\|_2 \geq \Omega(n)$, so that by Lemma 56 we have

$$\|\boldsymbol{H}\boldsymbol{x}\|_q \geq \|\boldsymbol{H}\boldsymbol{x}\|_2 \cdot n^{1/q-1/2} = \Omega(n^{1/2+1/q}).$$

For $p \in [1,2]$, take an identity matrix $\boldsymbol{I}_n$. For every $\boldsymbol{x} \in [-1,1]^n$ with $|\{i : x_i^2 = 1\}| \geq n/2$ we have $\|\boldsymbol{I}_n\boldsymbol{x}\|_q = \|\boldsymbol{x}\|_q \geq \Omega(n^{1/q})$, and the columns of $\boldsymbol{I}_n$ are certainly in B_p^m. □

2.4 Signed series in Lebesgue spaces

In this section, we generalize Theorem 51 for all prefixes:

Theorem 63. *Let $n \leq 2d$ and $2 \leq p \leq q \leq \infty$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$, there exists a polynomial-time computable partial coloring $\boldsymbol{x} \in [-1,1]^n$ with $|\{i : x_i^2 = 1\}| \geq n/2$ so that*

$$\max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_q \lesssim \sqrt{\min\left(p, \log\left(\tfrac{4d}{n}\right)\right)} \cdot n^{1/2-1/p+1/q}.$$

While it is known that the vector balancing constant is largest when $n = d$ up to a factor of two, the same is not known for all prefixes. On the other hand, for partial colorings, we can indeed reduce it to the case where $n \leq 2d$. This is because if $n > 2d$, we can split the vectors into $\lfloor n/2d \rfloor$ intervals I of size $2d$ (and at most one leftover block of size at most $2d$) and set $d \cdot \lfloor n/2d \rfloor$ linear constraints of the form $\sum_{i\in I} x_i a_{i,j} = 0$ for all $j \in [d]$. This induces a subspace of linear dimension so that by Theorem 7 it suffices to show a measure lower bound for this section of the corresponding discrepancy body.

By iteratively applying Theorem 63 we can obtain a full coloring at the expense of another factor of $\frac{1}{1/2-1/p+1/q}$, or $\log(n/d)$ if $n > 2d$:

Theorem 64. *Let $2 \leq p \leq q \leq \infty$ with $\{p,q\} \neq \{2,\infty\}$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$, there exist polynomial-time computable signs $\boldsymbol{x} \in \{-1,1\}^n$ so that, if $n \leq 2d$,*

$$\max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_q \lesssim \frac{\sqrt{\min\left(p, \log\left(\frac{4d}{n}\right)\right)}}{1/2 - 1/p + 1/q} \cdot n^{1/2-1/p+1/q}.$$

If instead $n > 2d$, the signs satisfy

$$\max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_q \lesssim \Big(\log(n/d) + \frac{1}{1/2 - 1/p + 1/q} \Big) \cdot d^{1/2-1/p+1/q}.$$

This improves upon the constructive bound of $\log(n)^{2.5}$ [72] to $\log(n)$ when $p = 2$ and $q = \infty$ (this is equivalent to taking $q = \log(d)$). When $p = q$ and $m = n$, we get the following corollary:

Corollary 65 (ℓ_p prefix version of Spencer's theorem)**.** *Let $2 \leq p \leq \infty$ and $n \in \mathbb{N}$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^n$, there exist polynomial-time computable signs $\boldsymbol{x} \in \{-1, 1\}^n$ so that*

$$\max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_p \lesssim \sqrt{n}.$$

By applying a theorem of Banaszczyk which relates signed series to the signed rearrangement problem, we obtain the following:

Corollary 66. *Let $2 \leq p \leq \infty$ and $n \in \mathbb{N}$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$, there exist polynomial-time computable signs $\boldsymbol{x} \in \{-1, 1\}^n$ and a permutation $\pi \in S_n$ so that*

$$\max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_{\pi(i)} \Big\|_p \lesssim \sqrt{d} \log\log d.$$

As before, Theorem 63 follows from a measure lower bound:

Theorem 67. *Let $n \leq 2d$ and $1 \leq p \leq q \leq \infty$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$,*

$$\gamma_n\Big(\Big\{ \boldsymbol{x} \in \mathbb{R}^n : \max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_q \leq \sqrt{\min\Big(p, \log\Big(\frac{4d}{n}\Big)\Big)} \cdot n^{\max(0, 1/2-1/p)+1/q} \Big\} \Big) \geq 2^{-O(n)}.$$

In order to show Theorem 67, we once again show the two special cases of $q \in \{p, \infty\}$ and bootstrap them into a general bound. First we address the simpler case $p = q$:

Lemma 68. *Let $p \geq 1$. Then for any $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$,*

$$\gamma_n\Big(\Big\{ \boldsymbol{x} \in \mathbb{R}^n : \max_{k\in[n]} \Big\| \sum_{i=1}^{k} x_i \boldsymbol{a}_i \Big\|_p \leq \sqrt{p} \cdot n^{\max(1/2, 1/p)} \Big\} \Big) \geq 2^{-O(n)}.$$

Proof. By Lemma 59 we know that, for some constant $C > 0$,

$$\mathop{\mathbb{E}}_{\boldsymbol{x}\sim N(\boldsymbol{0},\boldsymbol{I}_n)}\Big[\Big\|\sum_{i=1}^{n} x_i \boldsymbol{a}_i\Big\|_p\Big] \leq C\sqrt{p}\cdot n^{\max(1/2,1/p)},$$

for every $k \in [n]$. By Markov's inequality it follows that

$$\gamma_n\Big(\Big\{\boldsymbol{x}\in\mathbb{R}^n : \Big\|\sum_{i=1}^{k} x_i\boldsymbol{a}_i\Big\|_p \leq 2C\sqrt{p}\cdot n^{\max(1/2,1/p)}\Big\}\Big) \geq 1/2,$$

and Šidak's Lemma 5 gives

$$\gamma_n\Big(\bigcap_{k\in[n]}\Big\{\boldsymbol{x}\in\mathbb{R}^n : \Big\|\sum_{i=1}^{k} x_i\boldsymbol{a}_i\Big\|_p \leq 2C\sqrt{p}\cdot n^{\max(1/2,1/p)}\Big\}\Big) \geq 2^{-n},$$

so that the result follows by Lemma 16. □

Before we deal with the case $q = \infty$, we need a crucial lemma for a single row:

Lemma 69. *Let $\boldsymbol{a}\in\mathbb{R}^n$ so that $\|\boldsymbol{a}\|_2^2 \leq 1$. Then $\gamma_n(K_{\boldsymbol{a}}) \geq \exp(-C\cdot\exp(-\frac{1}{2\|\boldsymbol{a}\|_2^2}))$, where*

$$K_{\boldsymbol{a}} := \Big\{\boldsymbol{x}\in\mathbb{R}^n : \max_{k\in[n]}|\sum_{i=1}^{k} x_i a_i| \leq C_0\Big\},$$

for some universal constants $C, C_0 > 0$.

Proof. We may assume $n = 2^\ell$ is a power of two (otherwise we may append zeros). Let $y := \|\boldsymbol{a}\|_2^2$, and for a set $I \subseteq [n]$ let $s(I) := \sum_{i\in I} a_i^2$ so that $y = s([n])$. We also define $m(I)$ as the first index for which $s(I\cap\{i\in[n] : i \leq m(I)\}) > \frac{1}{2}s(I)$. Denote also $L(I) := I\cap\{i\in[n] : i < m(I)\}$ and $R(I) := I\cap\{i\in[n] : i > m(I)\}$ so that both $s(L(I)), s(R(I)) \leq \frac{1}{2}s(I)$.

We define $\ell+1$ families of intervals of $[n]$ as follows. The first family is $\mathcal{I}_0 := [n], m([n])$. For $i\in[\ell]$, we define $\mathcal{I}_i$ as $\bigcup_{I\in\mathcal{I}_{i-1}}\Big\{L(I), R(I), \{m(L(I))\}, \{m(R(I))\}\Big\}$. The key properties of these families are:

1. $|\mathcal{I}_i| = 2\cdot 4^{i-1} < 4^i$;

2. For any $I\in\mathcal{I}_i$ one has $s(I) \leq y/2^i$;

3. Any interval $I \subseteq [n]$ can be greedily expressed as a disjoint union of several sub-intervals in such a way that there are at most 2 sub-intervals from each $\mathcal{I}_i$, $i \in [\ell]$.

We claim for some constant $C_0 > 0$ it holds that

$$C_0 K_{\boldsymbol{a}} \supseteq \bigcap_{i=0}^{\ell} \bigcap_{I \in \mathcal{I}_i} \Big\{ \boldsymbol{x} \in \mathbb{R}^n : |\langle \boldsymbol{a}_I, \boldsymbol{x} \rangle| \leq \sqrt{\frac{1+4iy}{2^i}} \Big\} =: K'_{\boldsymbol{a}}.$$

Indeed, by the third property, for any $\boldsymbol{x} \in K'_{\boldsymbol{a}}$ and any $k \in [n]$ one has

$$\Big| \sum_{i=1}^{k} x_i a_i \Big| \leq 2 \cdot \sum_{i=0}^{\ell} \sqrt{\frac{1+4iy}{2^i}} \lesssim 1,$$

where we have also used that $y \leq 1$ by assumption.

By Šidak's Lemma 5, Lemma 15 and the second property we then have

$$\begin{aligned} \gamma_n(K'_{\boldsymbol{a}}) &\geq \prod_{i=0}^{\ell} \exp\Big(- \exp(-(1+4iy)/2^i/(2 \cdot y/2^i))\Big)^{|\mathcal{I}_\ell|} \\ &= \prod_{i=0}^{\ell} \exp(-\tfrac{1}{e^{2i}} \exp(-\tfrac{1}{2y}))^{|\mathcal{I}_\ell|}. \end{aligned}$$

Finally, by the first property we get

$$\gamma_n(K'_{\boldsymbol{a}}) \geq \exp\Big(- \sum_{i=0}^{\ell} (\tfrac{4}{e^2})^i \cdot \exp(-\tfrac{1}{2y})\Big) \geq \exp(-C \cdot \exp(-\tfrac{1}{2y})),$$

so that $\gamma_n(C_0 K_{\boldsymbol{a}}) \geq \exp(-C \cdot \exp(-\frac{1}{2y}))$. □

Lemma 70. *Let $n \leq 2d$ and $p \geq 2$. Then for any $\boldsymbol{A} \in \mathbb{R}^{d \times n}$ with columns $\boldsymbol{a}_1, \dots, \boldsymbol{a}_n \in B_p^d$ and rows $\boldsymbol{A}_1, \dots, \boldsymbol{A}_d \in \mathbb{R}^n$, the body $K := \{\boldsymbol{x} \in \mathbb{R}^n : \max_{k \in [n]} \| \sum_{i=1}^k x_i \boldsymbol{a}_i \|_\infty \leq \sqrt{p} \cdot n^{1/2-1/p}\}$ satisfies $\gamma_n(K \cap H) \geq 2^{-O(n)}$ for some subspace $H \subseteq \mathbb{R}^n$ with $\dim(H) \gtrsim n$.*

Proof. As before, we rescale the row vectors to $\boldsymbol{V}_j := (\sqrt{p} n^{\max(0,1/2-1/p)})^{-1} \boldsymbol{A}_j$ and abbreviate $y_j := \|\boldsymbol{V}_j\|_2^2$, so that by Eq. (2.1) we have $\sum_{j=1}^d y_j^{p/2} \leq n \cdot p^{-p/2}$. If any row satisfies $y_j > 1$, we may split it into maximal intervals of squared ℓ_2 norm at most 1, and impose a linear constraint in H that each maximal interval be orthogonal to the coloring; note that there are at most $n \cdot p^{-p/2}$ such constraints. We are left with dealing the case where all $y_j \leq 1$. Here we may apply Lemma 69 to conclude that the Gaussian measure of $K \cap H$ is at least

$$\exp(-C \cdot \sum_{j=1}^{d} \exp(-\tfrac{1}{2y_j})) \geq \exp(-C \cdot \sum_{j=1}^{d} y_j^{p/2} p^{p/2}) \geq \exp(-Cn). \qquad \square$$

Proof of Theorem 67. Now follows directly from bootstrapping Lemmas 68 and 70 as in the proof of Theorem 60. □

Proof of Theorem 64. We similarly iterate the partial colorings while observing that after each iteration we once again obtain a signed series problem. The only difference from Theorem 52 is that for $n > 2d$, we pay the extra factor of $\log(n/d)$ to bring the number of vectors down to the dimension. □

Proof of Corollary 65. Follows directly from Theorem 64. □

Proof of Corollary 66. Follows directly from Theorem 64 and a theorem of Banaszczyk [14] that the signed rearrangement constant is at most twice the signed series constant with $8d \log d$ many vectors. □

2.5 *Lower bounds*

Finally, we show that the bound in Theorem 52 is tight up to a factor of $\frac{1}{2} - \frac{1}{p} + \frac{1}{q}$:

Theorem 71. *Let $2 \leq p \leq q \leq \infty$ and $n \leq d \leq 2^n$. Then*

$$\mathsf{vb}_n(B_p^d, B_q^d) \gtrsim \sqrt{\min\left\{p, \log\left(\frac{2d}{n}\right)\right\}} \cdot n^{1/2-1/p+1/q}.$$

For the lower bound we will use the following reverse Chernoff bound from [88]:

Lemma 72. *Given independent random variables $x_1, \ldots, x_n \sim \{-1, 1\}$ and $\lambda \in [3, \sqrt{n}/2]$,*

$$\Pr[x_1 + \cdots + x_n \geq \lambda\sqrt{n}] \geq \exp(-9\lambda^2/2).$$

In this section we show that the vector balancing bounds in Theorem 52 are tight up to the factor of $\frac{1}{2} - \frac{1}{p} + \frac{1}{q}$:

Proof of Theorem 71. First, let us focus on the case when $d \le 2^p n$ so that $\log(2d/n) \le p$. We will also assume that $d \ge 8n$, since otherwise we can use instead $n' := \lfloor n/8 \rfloor$ and add $n - n'$ columns of zeros. Let $\boldsymbol{B}$ denote an $d \times n$ random matrix with i.i.d ± 1 entries. For any fixed $\boldsymbol{x} \in \{-1,1\}^n$, let $N_{\boldsymbol{x}}$ denote the number of rows $i \in [d]$ with $\langle \boldsymbol{B}_i, \boldsymbol{x} \rangle \ge \lambda\sqrt{n}$ for $\lambda := \sqrt{\frac{2}{9}\log\left(\frac{d}{2n}\right)}$. Since $\lambda \le \sqrt{n}/2$, by Lemma 72 we have

$$\Pr\left[\langle \boldsymbol{B}_i, \boldsymbol{x} \rangle \ge \lambda\sqrt{n}\right] \ge \frac{2n}{d},$$

so that $\mathbb{E}[N_{\boldsymbol{x}}] \ge 2n$. The standard Chernoff bound then gives $\Pr[N_{\boldsymbol{x}} \le (1-0.9)\cdot 2n] < 2^{-n}$, so that by the union bound there exists a matrix $\boldsymbol{B} \in \{-1,1\}^{d\times n}$ for which $N_{\boldsymbol{x}} \ge n/5$ for all $\boldsymbol{x} \in \{-1,1\}^n$. Thus for any such $\boldsymbol{x}$,

$$\|\boldsymbol{B}\boldsymbol{x}\|_q \ge (|N_{\boldsymbol{x}}| \cdot (\lambda\sqrt{n})^q)^{1/q} \gtrsim n^{1/2+1/q} \log\left(\frac{d}{2n}\right)^{1/2} \gtrsim n^{1/2+1/q} \log\left(\frac{2d}{n}\right)^{1/2},$$

where in the last step we used $d \ge 8n$. The matrix $\boldsymbol{A} := d^{-1/p}\boldsymbol{B}$ has columns in B_p^d and

$$\|\boldsymbol{A}\boldsymbol{x}\|_q \gtrsim \sqrt{\log\left(\frac{d}{2n}\right)} \cdot \frac{n^{1/2+1/q}}{d^{1/p}} \gtrsim \sqrt{\log\left(\frac{d}{2n}\right)} \cdot n^{1/2-1/p+1/q}$$

for all $\boldsymbol{x} \in \{-1,1\}^n$, as claimed.

For the case when $d \ge 2^p n$ we may use the same construction for $d' := \lfloor 2^p n \rfloor$ with $d - d'$ additional rows of zeros. □

2.6 Open problems

We conjecture that Theorem 52 can be improved to match Theorem 51:

Conjecture 6 ($\ell_p \to \ell_q$ version of Komlós conjecture)**.** *Given $n \le d$, $2 \le p \le q \le \infty$ and $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in B_p^d$, do there always exist signs $\boldsymbol{x} \in \{-1,1\}^n$ so that*

$$\Big\|\sum_{i=1}^n x_i \boldsymbol{a}_i\Big\|_q \lesssim \sqrt{\min\left(p, \log\left(\frac{2d}{n}\right)\right)} \cdot n^{1/2-1/p+1/q}.$$

Since Conjecture 6 is at least as hard as the Komlós conjecture, a more realistic goal would be to improve the full coloring of Theorem 52 by a factor of $(1/2 - 1/p + 1/q)^{-1/2}$ so as to match the best known bound of $O(\sqrt{\log n})$ for Komlós.

Recall that for a matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ and $1 \leq p \leq \infty$, the *Schatten-p norm* is defined as $\|\boldsymbol{A}\|_{S(p)} := (\sum_{i=1}^n \sigma_i(\boldsymbol{A})^p)^{1/p}$ where $\sigma_i(\boldsymbol{A}) \geq 0$ is the ith *singular value* of the matrix. In particular $\|\boldsymbol{A}\|_{S(\infty)}$ is the maximum singular value and $\|\boldsymbol{A}\|_{S(1)}$ is known as *Trace norm* or *Nuclear norm*. One might wonder whether Theorem 51 could be extended for *matrices* instead of vectors in the corresponding Schatten norms. In fact this is not possible: even for $p = 2$ and $q = \infty$, there exist n rank-one matrices $\boldsymbol{A}_i := \boldsymbol{v}_i\boldsymbol{v}_i^\top \in \mathbb{R}^{n\times n}$ with unit $\boldsymbol{v}_i$ for which any fractional coloring has discrepancy $\Omega(\sqrt{n})$ in the operator norm ([169], Section 3). It is still possible nevertheless that Corollary 53 extends in the following way:

Conjecture 7 (ℓ_p version of Matrix Spencer)**.** *Given $2 \leq p \leq \infty$ and symmetric $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_n \in \mathbb{R}^{n\times n}$ with Schatten-p norm at most 1, can we always find signs $\boldsymbol{x} \in \{-1, 1\}^n$ so that*

$$\Big\| \sum_{i=1}^n x_i \boldsymbol{A}_i \Big\|_{S(p)} \lesssim \sqrt{n}.$$

This is a more general form of the Matrix Spencer conjecture [171], and one can show a weaker bound of $O(\sqrt{pn})$ with random signs similar to Lemma 59. In fact, it is an open problem to show even a partial coloring for Conjecture 2. This would be implied by the following, which at least holds for diagonal matrices by the proof of Lemma 62:

Conjecture 8. *Given $1 \leq p \leq \infty$ and symmetric $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_n \in \mathbb{R}^{n\times n}$, can we show that*

$$K := \Big\{ \boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{A}_i \Big\|_{S(p)} \leq \Big\| \Big(\sum_{i=1}^n \boldsymbol{A}_i^2 \Big)^{1/2} \Big\|_{S(p)} \Big\}$$

satisfies $\gamma_n(K) \geq 2^{-O(n)}$?

2.7 *Proof of Lemma 59*

Proof of Lemma 59. By convexity of $z \mapsto |z|^p$, Jensen's inequality in $(*)$ and Khintchine's inequality in $(**)$ (Lemma 58) we have

$$\mathbb{E}\left[\Big\|\sum_{i=1}^{n} x_i \boldsymbol{a}_i\Big\|_p\right] \overset{(*)}{\leq} \mathbb{E}\left[\Big\|\sum_{i=1}^{n} x_i \boldsymbol{a}_i\Big\|_p^p\right]^{1/p}$$
$$= \Big(\sum_{j\in[d]} \mathbb{E}\left[\Big|\sum_{i\in[n]} x_i a_{ij}\Big|^p\right]\Big)^{1/p}$$
$$\overset{(**)}{\leq} C\sqrt{p}\cdot\Big(\sum_{j\in[d]}\Big(\sum_{i\in[n]} a_{ij}^2\Big)^{p/2}\Big)^{1/p}.$$

If $p \in [1,2]$, write $\boldsymbol{A}_j \in \mathbb{R}^n$ as $(\boldsymbol{A}_j)_i := a_{ij}$. Then by Lemma 56,

$$\Big(\sum_{j\in[d]}\Big(\sum_{i\in[n]} a_{ij}^2\Big)^{p/2}\Big)^{1/p} = \Big(\sum_{j\in[d]}\|\boldsymbol{A}_j\|_2^p\Big)^{1/p} \leq \Big(\sum_{j\in[d]}\|\boldsymbol{A}_j\|_p^p\Big)^{1/p} = \Big(\sum_{i\in[n]}\|\boldsymbol{a}_i\|_p^p\Big)^{1/p} \leq n^{1/p}.$$

Now suppose that $p \geq 2$. Define $(\boldsymbol{a}_i)^2 \in \mathbb{R}^d$ to be the vector with jth coordinate a_{ij}^2. Since $\|\cdot\|_{p/2}$ is a norm, we can use the triangle inequality to get

$$\Big(\sum_{j\in[d]}\Big(\sum_{i\in[n]} a_{ij}^2\Big)^{p/2}\Big)^{1/p} = \Big\|\sum_{i\in[n]}(\boldsymbol{a}_i)^2\Big\|_{p/2}^{1/2} \leq \Big(\sum_{i\in[n]}\|(\boldsymbol{a}_i)^2\|_{p/2}\Big)^{1/2} = \Big(\sum_{i\in[n]}\|\boldsymbol{a}_i\|_p^2\Big)^{1/2} \leq n^{1/2}.$$

Either way, we conclude that $\mathbb{E}[\|\sum_{i=1}^n x_i\boldsymbol{a}_i\|_p] \leq O(\sqrt{p}\cdot n^{\max(1/2,1/p)})$, as desired. $\square$

Remark 4. *A similar approach gives an alternate proof of Prop. 25 in [30], which states that a $r := O(\sqrt{p}\cdot n^{1/p})$ scaling of an n-dimensional section H of B_p^d has Gaussian measure $\gamma_H(H \cap rB_p^d) \geq 1/2$ for $p \geq 2$. Indeed, by Markov's inequality, it suffices to note that given an orthonormal basis $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n$ of H we have*

$$\mathbb{E}\left[\Big\|\sum_{i=1}^{n} x_i \boldsymbol{a}_i\Big\|_p\right] \leq C\sqrt{p}\cdot\Big(\sum_{j\in[d]}\Big(\sum_{i\in[n]} a_{ij}^2\Big)^{p/2}\Big)^{1/p} \leq C\sqrt{p}\cdot n^{1/p},$$

where the last inequality follows from convexity of $z \mapsto z^{p/2}$ and from the fact that the m terms $\sum_{i\in[n]} a_{ij}^2$ sum to n and are at most 1 by orthonormality.

Chapter 3

THE VECTOR BALANCING CONSTANT OF ZONOTOPES

This chapter is based on a joint paper with Rainie Bozzai and Thomas Rothvoss [140].

3.1 Introduction

Discrepancy theory is a subfield of combinatorics where one is given a set system $(X, \mathcal{F})$ with a ground set X and a family of sets $\mathcal{F} \subseteq 2^X$, and the goal is to find the coloring that minimizes the maximum imbalance, i.e.

$$\operatorname{disc}(\mathcal{F}) = \min_{x \in \{-1,1\}^X} \max_{S \in \mathcal{F}} \Big| \sum_{j \in S} x_j \Big|.$$

A slightly more general linear-algebraic view is that one is given a matrix $A \in [-1,1]^{d \times n}$ and its discrepancy is defined as $\min_{x \in \{-1,1\}^n} \|Ax\|_\infty$. The best known result in this area is certainly Spencer's Theorem [157] which states that for any $n \leq d$ one has $\operatorname{disc}(A) \leq O(\sqrt{n \log(\frac{2d}{n})})$. The challenging aspect of that Theorem is that — say for $n = d$ — a uniform random coloring $x \sim \{-1,1\}^n$ will only give a $\Theta(\sqrt{n \log n})$ bound. Instead, Spencer [157] applied the *partial coloring method* which had been first used by Beck [32].

The original proofs of the partial coloring method are based on the *pigeonhole principle* and are non-constructive. The first polynomial time algorithm to actually find the coloring guaranteed by Spencer [157] is due to Bansal [16], followed by a sequence of algorithms [103, 146, 100, 61] that either work in more general settings or are simpler.

Discrepancy theory is an extensively studied topic with many applications in mathematics and computer science. To give two concrete examples, Nikolov, Talwar and Zhang [126] showed a connection between differential privacy and hereditary discrepancy, and the best known approximation algorithm for Bin Packing uses a discrepancy-based rounding [74].

Other applications can be found in data structure lower bounds, communication complexity and pseudorandomness; we refer to the book of Chazelle [45] for a more detailed account. The seminal result of Batson, Spielman and Srivastava [31] on the existence of linear-size spectral sparsifiers for graphs can also be interpreted as a discrepancy-theoretic result, see [138] for details.

Spencer's Theorem [157] can be rephrased as $\mathsf{vb}(B_\infty^d, B_\infty^d) = \Theta(\sqrt{d})$ and as $\mathsf{vb}_n(B_\infty^d, B_\infty^d) = \Theta(\sqrt{n\log(\frac{2d}{n})})$ for $n \leq d$. Here we denote B_p^d as the d-dimensional unit ball of the norm $\|\cdot\|_p$. Moreover for a Euclidean ball one can easily prove that $\mathsf{vb}(B_2^d, B_2^d) = \Theta(\sqrt{d})$ and for the ℓ_1-ball we have $\mathsf{vb}(B_1^d, B_1^d) = \Theta(d)$.

While Spencer's Theorem itself is tight, at least three candidate generalizations have been suggested in the literature — all three are unsolved so far.

The Beck-Fiala Conjecture. Suppose we have a set system $(X, \mathcal{F})$ in which every element is in at most t sets. Beck and Fiala [33] proved using a linear-algebraic argument that in this case the discrepancy is bounded by $2t$ and they state the conjecture that the correct dependence should be $O(\sqrt{t})$. The same proof of [33] also shows that $\mathsf{vb}(B_1^d, B_\infty^d) \leq 2$. However, the Beck-Fiala Conjecture is wide open and the best known bounds are $O(\sqrt{t\log n})$ [13, 19] and $2t - \log^*(t)$ [40]. In fact, Komlós Conjecture of $\mathsf{vb}(B_2^d, B_\infty^d) \leq O(1)$ is even more general; here the best known bound is $\mathsf{vb}(B_2^d, B_\infty^d) \leq O(\sqrt{\log(d)})$ [13].

The Matrix Spencer Conjecture. A conjecture popularized by Zouzias [171] and Meka [113] claims that for any symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{n\times n}$ with all eigenvalues in $[-1, 1]$, there are signs $x \in \{-1, 1\}^n$ so that the maximum singular value of $\sum_{i=1}^n x_i A_i$ is at most $O(\sqrt{n})$. Using standard matrix concentration bounds, one can prove that a random coloring attains a value of at most $O(\sqrt{n\log n})$. Moreover, one can prove the conjectured upper bound of $O(\sqrt{n})$ under the additional assumption that the matrices are block-diagonal with constant size blocks [52], or have rank $O(\sqrt{n})$ [76]. Based on recent progress on matrix concentration, it is possible to obtain the same under the weaker condition that they have rank at most $\frac{n}{\log^3(n)}$ [21].

The vector balancing constant of zonotopes. A *zonotope* is defined as the linear image of a cube. If $A \in \mathbb{R}^{m\times d}$ is a matrix with $m \geq d$, we can write a d-dimensional zonotope in the form $K = \{\sum_{i=1}^m y_i A_i \mid y \in [-1,1]^m\} = A^\top B_\infty^m \subseteq \mathbb{R}^d$. Note that m is the *number of segments* of the zonotope. The cube B_∞^d is trivially a zonotope, and it is known that for every $p \geq 2$, the ball B_p^n is the limit of a sequence of zonotopes, called a *zonoid* [36]. Schechtman [152] raised the question whether it is true that for any zonotope $K \subseteq \mathbb{R}^d$ one has $\mathsf{vb}(K,K) \lesssim \sqrt{d}$ where we write $A \lesssim B$ if $A \leq C \cdot B$ for a universal constant $C > 0$. The best known bound of $\mathsf{vb}(K,K) \lesssim \sqrt{d \log\log d}$ is a direct consequence of Spencer's theorem and the fact that zonotopes can be *sparsified* up to a constant factor with only $O(d\log d)$ segments [162]. An affirmative answer to Schechtman's question would follow from an $O(d)$ bound, or equivalently whether an ℓ_1-analogue of [31] is true. We defer to Section 3.6 for details.

3.1.1 *Our contributions*

Our main result is an almost-proof of Schechtman's conjecture (falling short only by a $\log\log\log d$ term).

Theorem 73. *For any zonotope $K \subseteq \mathbb{R}^d$ one has $\mathsf{vb}(K,K) \lesssim \sqrt{d}\log\log\log d$. Moreover, for any $v_1,\dots,v_n \in K$ one can find in randomized polynomial time a coloring $x \in \{-1,1\}^n$ with $\|\sum_{i=1}^n x_i v_i\|_K \lesssim \sqrt{d}\log\log\log d$.*

The claim is invariant under linear transformations to K and so it will be useful to place K in a normalized position. For this sake, we make the following definition:

Definition 74. *A matrix $A \in \mathbb{R}^{m\times d}$ is called* approximately regular *if the following holds:*

(*i*) *The columns $A^1,\dots,A^d$ are orthonormal.*
(*ii*) *The rows satisfy $\|A_i\|_2 \leq 2\sqrt{\frac{d}{m}}$ for all $i = 1,\dots,m$.*

Then we call a zonotope $K \subseteq \mathbb{R}^d$ *normalized* if there exists a matrix $A \in \mathbb{R}^{m \times d}$ that is approximately regular so that $K = \sqrt{\frac{d}{m}} A^\top B_\infty^m$. We choose the scaling so that any cube B_∞^d is indeed normalized and zonotopes with any number of segments are comparable to B_∞^d in terms of volume and radius.

Our main technical contribution is a tight lower bound for the Gaussian measure of sections of any normalized zonotope.

Theorem 75. *For any normalized zonotope $K \subseteq \mathbb{R}^d$, any subspace $H \subseteq \mathbb{R}^d$ with $n := \dim(H)$ and any $t \geq 1$, one has $\gamma_H(t \cdot C \cdot K \cap H) \geq \exp(-e^{-t^2/2} \cdot n)$ where $C > 0$ is a universal constant.*

In order to prove Theorem 75, we show that a normalized zonotope can be decomposed into $\Theta(\frac{m}{d})$ many smaller zonotopes with $\Theta(d)$ many segments each. This decomposition requires an iterative application of the Kadison-Singer theorem by Marcus, Spielman and Srivastava [106]. Then we prove the statement of Theorem 75 for such simpler zonotopes and derive the lower bound on $\gamma_H(t \cdot C \cdot K \cap H)$ by using log-concavity of the Gaussian measure.

We can also use Theorem 75 to show how to balance vectors between different normalized zonotopes:

Theorem 76. *For any normalized zonotopes $K, Q \subseteq \mathbb{R}^d$ one has $\mathsf{vb}(K, Q) \lesssim \sqrt{d \log d}$. Moreover, for any $v_1, \ldots, v_n \in K$ one can find in randomized polynomial time a coloring $x \in \{-1, 1\}^n$ such that $\| \sum_{i=1}^n x_i v_i \|_Q \lesssim \sqrt{d \cdot \log \min\{d, n\}}$.*

3.2 *Preliminaries*

We review a few facts that we rely on later.

Probability. By γ_n we denote the (*standard*) *Gaussian density* $\frac{1}{(2\pi)^{n/2}} e^{-\|x\|_2^2/2}$. For the corresponding distribution we will write $N(0, I_n)$. For a subspace $F \subseteq \mathbb{R}^n$ we write $I_F \in \mathbb{R}^{n \times n}$ as the identity on the subspace; in particular $I_F = \sum_{i=1}^{\dim(F)} u_i u_i^T$ where $u_1, \ldots, u_{\dim(F)}$ is any orthonormal basis of F.

Discrepancy theory. First we give a full statement of Spencer's theorem that we mentioned earlier:

Theorem 77 (Spencer's Theorem [157, 103])**.** *For any $A \in [-1,1]^{m \times n}$ with $m \geq n$ there are polynomial time computable signs $x \in \{-1,1\}^n$ so that $\|Ax\|_\infty \lesssim \sqrt{n \log(\frac{2m}{n})}$. More generally, for any shift $x_0 \in [-1,1]^n$, there is a polynomial time computable $x \in \mathbb{R}^n$ so that $x + x_0 \in \{-1,1\}^n$ and $\|A(x+x_0)\|_\infty \lesssim \sqrt{n \log(\frac{2m}{n})}$.*

To be exact, the first algorithm giving a bound of $O(\sqrt{n}\log(\frac{2m}{n}))$ is due to Bansal [16] and the tight algorithmic bound is due to Lovett and Meka [103].

We say that a vector $x \in \mathbb{R}^n$ is a *good partial coloring* if $x \in [-1,1]^n$ with $|\{j \in [n] : x_j \in \{-1,1\}\}| \geq n/2$. We will need a connection between good partial colorings and Gaussian measure lower bounds.

Theorem 78 (special case of Theorem 7)**.** *For any $\alpha > 0$, there is a constant $c := c(\alpha) > 0$ and a randomized polynomial time algorithm that for a symmetric convex body $K \subseteq \mathbb{R}^n$, a $2n/3$-dimensional subspace $F \subseteq \mathbb{R}^n$ with $\gamma_F(K \cap F) \geq e^{-\alpha n}$ and a shift $y \in (-1,1)^n$, finds $x \in c \cdot K \cap F$ so that $x + y$ is a good partial coloring.*

For many decades, the *Kadison-Singer problem* was an open question in operator theory. It was finally resolved in 2015:

Theorem 79 (Marcus, Spielman, Srivastava [106])**.** *Let $v_1, \ldots, v_m \in \mathbb{R}^d$ so that $\sum_{i=1}^m v_i v_i^\top = I_d$ and let $\varepsilon > 0$ so that $\|v_i\|_2^2 \leq \varepsilon$ for all $i \in [m]$. Then there is a partition $[m] = S_1 \dot\cup S_2$ so that for both $j \in \{1,2\}$ one has*

$$\Big\| \sum_{i \in S_j} v_i v_i^\top - \frac{1}{2} I_d \Big\|_{\text{op}} \leq 3\sqrt{\varepsilon}$$

In the definition of $\text{vb}(K,Q)$, there is no upper bound on the number of vectors to be balanced. But it is well-known that up to a constant factor, the worst-case is attained for d many vectors. Let

$$\text{vb}_n(K,Q) := \inf\Big\{ r \geq 0 \mid \forall u_1, \ldots, u_n \in K : \exists x \in \{-1,1\}^n : \sum_{j=1}^n x_j v_j \in rQ \Big\}$$

be the vector balancing variant with n vectors, so that $\mathsf{vb}(K,Q) := \sup_{n \in \mathbb{N}} \mathsf{vb}_n(K,Q)$.

Zonotopes. A substantial amount of work in the literature has been done on the question of how one can sparsify an arbitrary zonotope with another zonotope that has fewer segments, while losing only a constant factor approximation. The first bound of $O(d^2)$ [151] was improved to $O(d\log^3 d)$ [36]. We highlight the current best known bound:

Theorem 80 (Talagrand [162])**.** *For any zonotope $K \subseteq \mathbb{R}^d$ and $0 < \varepsilon \leq \frac{1}{2}$, there is a zonotope Q with at most $O(\frac{d}{\varepsilon^2}\log d)$ segments so that $Q \subseteq K \subseteq (1+\varepsilon)Q$.*

We refer to the approach of Cohen and Peng [47] for an elementary exposition of the $O(d\log d)$ bound.

Finally, we justify why it suffices to consider normalized zonotopes:

Lemma 81. *For any full-dimensional zonotope $K = A^\top B_\infty^m \subseteq \mathbb{R}^d$, there is a normalized zonotope $\tilde{K}$ and an invertible linear map T so that $\frac{4}{5}\tilde{K} \subseteq T(K) \subseteq \tilde{K}$. In particular, $\frac{4}{5}\mathsf{vb}(\tilde{K},\tilde{K}) \leq \mathsf{vb}(K,K) \leq \frac{5}{4}\mathsf{vb}(\tilde{K},\tilde{K})$.*

We show the argument in Section 3.7.

Lemma 82. *Any normalized zonotope $K \subseteq \mathbb{R}^d$ satisfies $K \subseteq \sqrt{d}B_2^d$.*

Proof. We write $K = \sqrt{\frac{d}{m}}A^\top B_\infty^m$ where $A \in \mathbb{R}^{m\times d}$. Note that $A^\top A = I_d$ by orthonormality of the columns of A and so $\|A\|_{\mathrm{op}} = \|A^\top A\|_{\mathrm{op}}^{1/2} = 1$. By definition, for any $x \in K$ there is a $y \in B_\infty^m$ with $x = \sqrt{\frac{d}{m}}A^\top y$, so that

$$\|x\|_2 = \sqrt{\frac{d}{m}}\|A^\top y\|_2 \leq \sqrt{\frac{d}{m}}\|A^\top\|_{\mathrm{op}} \cdot \|y\|_2 \leq \sqrt{d}. \qquad \square$$

3.3 Sections of normalized zonotopes

In this section we prove Theorem 75, showing that all sections of normalized zonotopes are large. In the most basic form where $K = B_\infty^d$ is a cube and $t = 1$, the statement is similar to a result of Vaaler [165] who proved that $\mathrm{Vol}_H(K \cap H) \geq 2^n$ for any n-dimensional subspace $H \subseteq \mathbb{R}^d$; though the geometry of a zonotope is more complex and the proof strategy is rather different.

3.3.1 A first direct lower bound

We begin with a simple estimate on the Gaussian measure of the section of a zonotope where we drop the scalar of $\sqrt{\frac{d}{m}}$. Hence this bound will be tight if the number of segments is close to d but rather loose otherwise. We denote Π_H as the orthogonal projection into a subspace H.

Lemma 83. *Let $K := A^\top B_\infty^m \subseteq \mathbb{R}^d$ be a zonotope where $A \in \mathbb{R}^{m\times d}$ is a matrix with orthonormal columns. Then for any subspace $H \subseteq \mathbb{R}^d$ with $n := \dim(H)$ and any $t \geq 1$ one has $\gamma_H(t \cdot K \cap H) \geq \exp(-e^{-t^2/2} \cdot n)$.*

Proof. Let $U \in \mathbb{R}^{d\times n}$ be a matrix with orthonormal columns $U^1, \ldots, U^n$ spanning H. Then if we draw $y \sim N(0, I_n)$, Uy is indeed a standard Gaussian in the subspace H. By assumption, $\sum_{i=1}^m A_i A_i^\top = I_d$, and this can be used to write any outcome of the random process as

$$Uy = \sum_{j=1}^n y_j I_d U^j = \sum_{i=1}^m A_i \sum_{j=1}^n y_j \left\langle A_i, U^j \right\rangle = \sum_{i=1}^m A_i \left\langle y, U^\top A_i \right\rangle. \tag{3.1}$$

Here one should think of $U^\top A_i \in \mathbb{R}^n$ as the coordinates of $\Pi_H(A_i)$ in terms of the basis U of H. From the expression in (3.1) we can draw the following conclusion:
Claim I. *For any $y \in \mathbb{R}^n$ and $s > 0$ one has $(|\left\langle y, U^\top A_i \right\rangle| \leq s\ \forall i \in [m]) \Rightarrow Uy \in sK$.*
Then Claim I gives a simple sufficient (but in general not necessary) condition for Uy to lie in the zonotope K. Next, we can see that

$$\sum_{i=1}^m \|U^\top A_i\|_2^2 = \sum_{i=1}^m \mathrm{Tr}\left[UU^\top A_i A_i^\top\right] = \mathrm{Tr}[UU^\top] = n$$

Then we can use Claim I and the inequality of Šidák-Khatri to lower bound the Gaussian measure by

$$\begin{aligned}\gamma_H(t\cdot K\cap H) &= \Pr_{y\sim N(0,I_n)}[Uy\in t\cdot K]\\ &\geq \Pr_{y\sim N(0,I_n)}\left[|\left\langle U^\top A_i,y\right\rangle|\leq t\ \ \forall i\in[m]\right]\\ &\overset{\text{Lem 5}}{\geq} \prod_{i=1}^m \Pr_{y\sim N(0,I_n)}\left[|\left\langle U^\top A_i,y\right\rangle|\leq t\right]\\ &\overset{\text{Lem 15}}{\geq} \prod_{i=1}^m \exp\big(-e^{-t^2/2}\|U^\top A_i\|_2^2\big)\\ &= \exp\Big(-e^{-t^2/2}\sum_{i=1}^m\|U^\top A_i\|_2^2\Big)=\exp\big(-e^{t^2/2}n\big)\end{aligned}$$

Here we have used that $\|U^\top A_i\|_2\leq\|A_i\|_2\leq 1$ which follows by the orthonormality of the columns of A. $\square$

It is somewhat unfortunate that Claim I shown above requires that $\sum_{i=1}^m A_iA_i^\top$ is *exactly* the identity and an approximation is not enough. But we can fix this by a rescaling argument:

Lemma 84. *Let $K=A^\top B_\infty^m\subseteq\mathbb{R}^d$ be a zonotope where $A\in\mathbb{R}^{m\times d}$ is a matrix so that $\sum_{i=1}^m A_iA_i^\top\succeq\alpha I_d$ for some $\alpha>0$. Then for any n-dimensional subspace $H\subseteq\mathbb{R}^d$ and any $t\geq 1$ one has $\gamma_H(\frac{t}{\sqrt{\alpha}}\cdot K\cap H)\geq\exp\big(-e^{-t^2/2}\cdot n\big)$.*

Proof. Scaling K by $\frac{1}{\sqrt{\alpha}}$ is equivalent to scaling $\sum_{i=1}^m A_iA_i^\top$ by $\frac{1}{\alpha}$, hence we may assume that indeed $\alpha=1$. Abbreviate $M:=\sum_{i=1}^m A_iA_i^\top\succeq I_d$ which is a symmetric positive definite matrix. Consider the matrix $\tilde{A}\in\mathbb{R}^{m\times d}$ with rescaled rows $\tilde{A}_i:=M^{-1/2}A_i$, so that $\sum_{i=1}^m\tilde{A}_i\tilde{A}_i^\top=I_d$. Let $\tilde{K}:=\tilde{A}^\top B_\infty^m=M^{-1/2}(K)$ and $\tilde{H}:=M^{-1/2}(H)$ be the rescaled zonotope and subspace. Let $U^1,\dots,U^n$ be an orthonormal basis of H. Then with $\tilde{U}=M^{-1/2}U$, $\tilde{U}^1,\dots,\tilde{U}^n$ will be the basis of $\tilde{H}$, but it will not be orthogonal in general. However, for $y\sim N(0,I_n)$ one has $\mathrm{Cov}(\tilde{U}y)=\tilde{U}\tilde{U}^\top=M^{-1/2}UU^\top M^{-1/2}\preceq I_{\tilde{H}}$. Then

$$\Pr_{y\sim N(0,I_d)}[Uy\in tK]=\Pr_{y\sim N(0,I_d)}[\tilde{U}y\in t\tilde{K}]\overset{\text{Lem 21}}{\geq}\Pr_{y\sim N(0,I_{\tilde{H}})}[y\in t\tilde{K}]\overset{\text{Lem 83}}{\geq}\exp\big(-e^{-t^2/2}n\big).\quad\square$$

3.3.2 Decomposition of normalized zonotopes

The next step in our proof strategy is to decompose the rows of an approximately regular matrix $A \in \mathbb{R}^{m\times d}$ into $\Theta(\frac{m}{d})$ many blocks $J \subseteq [m]$ so that $\sum_{i\in J} A_i A_i^\top \succeq \Omega(\frac{d}{m}) \cdot I_d$. For this purpose, we formulate a slight variant of Theorem 79.

Lemma 85. *Let $v_1, \dots, v_m \in \mathbb{R}^d$ be vectors with $\sum_{i=1}^m v_i v_i^\top \succeq L \cdot I_d$ for some $L > 0$ and let $\varepsilon := \max_{i=1,\dots,m} \|v_i\|_2^2$. Then there is a partition $[m] = S_1 \dot\cup S_2$ so that*

$$\sum_{i\in S_j} v_i v_i^\top \succeq \Big(\frac{L}{2} - 3\sqrt{L\varepsilon}\Big) I_d \quad \forall j \in \{1,2\}$$

Proof. Abbreviate $M := \sum_{i=1}^m v_i v_i^\top$ which is a PSD matrix with $M \succeq L \cdot I_d$. Define $v_i' := M^{-1/2} v_i$. Then $\sum_{i=1}^m v_i'(v_i')^\top = M^{-1/2}\Big(\sum_{i=1}^m v_i v_i^\top\Big) M^{-1/2} = I_d$. We set $\varepsilon' := \frac{\varepsilon}{L}$ and verify that for all i one has $\|v_i'\|_2^2 = v_i^\top M^{-1} v_i \leq v_i^\top(\frac{1}{L} I_d) v_i = \frac{\|v_i\|_2^2}{L} \leq \varepsilon'$. Then we apply Theorem 79 to the vectors $\{v_i'\}_{i\in[m]}$ and obtain a partition $[m] = S_1 \dot\cup S_2$ so that for $j \in \{1,2\}$ one has

$$M^{-1/2}\Big(\sum_{i\in S_j} v_i v_i^\top\Big) M^{-1/2} = \sum_{i\in S_j} v_i'(v_i')^\top \overset{\text{Thm 79}}{\succeq} \Big(\frac{1}{2} - 3\sqrt{\varepsilon/L}\Big) I_d,$$

and using the fact that $A \succeq B \implies M^{1/2} A M^{1/2} \succeq M^{1/2} B M^{1/2}$, we conclude

$$\sum_{i\in S_j} v_i v_i^\top \succeq \Big(\frac{1}{2} - 3\sqrt{\varepsilon/L}\Big) M^{1/2} I_d M^{1/2} \succeq \Big(\frac{L}{2} - 3\sqrt{L\varepsilon}\Big) I_d. \qquad \square$$

Now to the main lemma of this section where we decompose an approximately regular matrix by iteratively applying Lemma 85.

Lemma 86. *There is a universal constant $C > 0$ so that the following holds. Let $A \in \mathbb{R}^{m\times d}$ be an approximately regular matrix. Then there are disjoint subsets $J_1 \dot\cup \cdots \dot\cup J_k \subseteq [m]$ with $k \geq \frac{m}{Cd}$ and $|J_\ell| \leq Cd$ and $\sum_{i\in J_\ell} A_i A_i^\top \succeq \frac{1}{Ck} I_d$ for all $\ell \in [k]$.*

Proof. If $\frac{m}{d} \leq C$ we may set $k = 1$ and $J_1 = [m]$, so assume $m \geq Cd$. Set $\varepsilon := 4\frac{d}{m}$ so that $\|A_i\|_2^2 \leq \varepsilon$ for all $i \in [m]$. Let $t \in \mathbb{N}$ be a parameter that we choose later. For $s \in \{0, \dots, t\}$ we will obtain partitions $\mathcal{P}_s$ of the row indices starting with $\mathcal{P}_0 := \{[m]\}$ so that $\mathcal{P}_{s+1}$ is a refinement of $\mathcal{P}_s$ and moreover $|\mathcal{P}_s| = 2^s$. More precisely, in each iteration $s \in \{0, \ldots, t-1\}$ and for each $S \in \mathcal{P}_s$, we apply Lemma 85 to the vectors $\{A_i\}_{i \in S}$; if $S = S_1 \dot\cup S_2$ is the obtained partition, then we add $\{S_1, S_2\}$ to $\mathcal{P}_{s+1}$. We first analyze the corresponding eigenvalue lower bound. Define $L_s := 2^{-s} - 15\sqrt{2^{-s}\varepsilon}$.

Claim. *If $2^t \leq \frac{m}{Cd}$ for a large enough constant $C > 0$, then for all $s \in \{0, \ldots, t\}$ one has $\sum_{i \in S} A_i A_i^\top \succeq L_s I_d$ for all $S \in \mathcal{P}_s$.*

Proof of Claim. Clearly $L_s \leq 2^{-s}$ all $s \geq 0$. We will prove the claim by induction on s. For $s = 0$ one has $\mathcal{P}_0 = \{[m]\}$ and the claim is true as $L_0 \leq 1$. Now consider an iteration $s \in \{0, \ldots, t-1\}$ and suppose $S \in \mathcal{P}_s$ is split into $S = S_1 \dot\cup S_2$. Then $\sum_{i \in S_j} A_i A_i^\top \succeq (\frac{L_s}{2} - 3\sqrt{L_s \varepsilon}) I_d$ for both $j \in \{1, 2\}$. This is at least L_{s+1} as:

$$\frac{L_s}{2} - 3\sqrt{L_s\varepsilon} \overset{L_s \leq 2^{-s}}{\geq} \frac{L_s}{2} - 3\sqrt{2^{-s}\varepsilon} \geq 2^{-(s+1)} - \frac{15}{2}\sqrt{2^{-s}\varepsilon} - 3\sqrt{2^{-s}\varepsilon} \geq 2^{-(s+1)} - 15\sqrt{2^{-(s+1)}\varepsilon}.$$

Here we use $15/2 + 3 \leq 15\sqrt{2^{-1}}$. This shows the claim. $\square$

For a large enough constant C, we pick $t \in \mathbb{N}$ so that $\frac{m}{2Cd} \leq 2^t \leq \frac{m}{Cd}$. Then $L_t \geq \frac{Cd}{m} - 15\sqrt{\frac{2Cd}{m} \cdot 4\frac{d}{m}} = \frac{d}{m} \cdot (C - 15\sqrt{8C}) \geq \frac{C}{2} \cdot \frac{d}{m}$ for C large enough. Moreover we know that $\mathbb{E}_{S \sim \mathcal{P}_t}[|S|] = \frac{m}{2^t} \leq 2Cd$. Then by Markov's inequality at least half the sets $S \in \mathcal{P}_t$ have at most $4Cd$ indices. Those sets will satisfy the statement. $\square$

3.3.3 *Proof of Theorem 75*

Next we prove our main technical result, Theorem 75. Recall that a measure μ on $\mathbb{R}^d$ is called *log-concave* if for all compact subsets $S, T \subseteq \mathbb{R}^d$ and $0 \leq \lambda \leq 1$ one has

$$\mu(\lambda S + (1-\lambda)T) \geq \mu(S)^\lambda \cdot \mu(T)^{1-\lambda}$$

By induction one can verify that for any compact subsets $S_1, \ldots, S_k \subseteq \mathbb{R}^d$ and $\lambda_1, \ldots, \lambda_k \geq 0$ with $\sum_{i=1}^k \lambda_i = 1$ we have $\mu(\lambda_1 S_1 + \cdots + \lambda_k S_k) \geq \prod_{\ell=1}^k \mu(S_\ell)^{\lambda_\ell}$. Also recall that the Gaussian measure γ_d is indeed log-concave, see e.g. [7]. For a matrix $A \in \mathbb{R}^{m \times d}$ and indices $J \subseteq [m]$ we denote $A_J \in \mathbb{R}^{|J| \times d}$ as the submatrix of A with rows in J.

Proof of Theorem 75. Let $K \subseteq \mathbb{R}^d$ be a normalized zonotope and let $H \subseteq \mathbb{R}^d$ be a subspace with dimension n. Then we can write $K = \sqrt{\frac{d}{m}} A^\top B_\infty^m$ where $A \in \mathbb{R}^{m \times d}$ is approximately regular. We use Lemma 86 to obtain disjoint subsets $J_1 \dot\cup \cdots \dot\cup J_k \subseteq [m]$ with $k \geq \frac{m}{Cd}$ so that $\sum_{i \in J_\ell} A_i A_i^\top \succeq \frac{d}{Cm} I_d$ where $C > 0$ is a constant. Consider the zonotope $K_\ell := \sqrt{\frac{d}{m}} A_{J_\ell}^\top B_\infty^{|J_\ell|}$ generated by the rows with indices in J_ℓ. Then we have $K_1 + \ldots + K_k \subseteq K$ and $(K_1 \cap H) + \ldots + (K_k \cap H) \subseteq K \cap H$. Note that for each $\ell \in [k]$ we have $kK_\ell \supseteq \sqrt{\frac{k}{C}} A_{J_\ell}^\top B_\infty^{|J_\ell|}$, so that $\sum_{i \in J_\ell} (\sqrt{\frac{k}{C}} A_i)(\sqrt{\frac{k}{C}} A_i)^\top \succeq \frac{k}{C} \cdot \frac{d}{Cm} I_d \succeq \frac{1}{C^3} I_d$. Then applying Lemma 84 with $\alpha := \frac{1}{C^3}$ we have

$$\gamma_H\big(tC^{3/2} kK_\ell \cap H\big) \geq \exp\big(-e^{t^2/2} \cdot n\big)$$

for all $t \geq 1$. Finally, using log-concavity of the Gaussian measure we obtain

$$\begin{aligned}\gamma_H\big(tC^{3/2} K \cap H\big) &\geq \gamma_H\big((tC^{3/2} K_1 \cap H) + \ldots + (tC^{3/2} K_k \cap H)\big) \\ &\geq \prod_{\ell=1}^k \gamma_H\big(tC^{3/2} \cdot kK_\ell \cap H\big)^{1/k} \geq \exp\big(-e^{-t^2/2} \cdot n\big).\end{aligned}$$ □

3.4 *The vector balancing constant* $\mathrm{vb}(K, K)$

Next, we show how to translate measure lower bounds for sections into an improved bounds on the vector balancing constant.

3.4.1 *Tight partial colorings for zonotopes*

First we prove a generalization of the constant discrepancy partial coloring for the Komlós setting:

Lemma 87. *Let $v_1, \ldots, v_n \in B_2^d$ and let $K \subseteq \mathbb{R}^d$ be a symmetric convex body with $\gamma_H(K \cap H) \geq e^{-\alpha n}$ for some $\alpha > 0$ where $H = \mathrm{span}\{v_1, \ldots, v_n\}$. Then there is a randomized polynomial time algorithm that given a shift $y \in (-1,1)^n$ finds a good partial coloring $x + y \in [-1,1]^n$ with $\sum_{j=1}^n x_j v_j \in cK$ where $c := c(\alpha)$ is a constant.*

Proof. Let $Z \sim \sum_{j=1}^n z_j v_j$ where $z_i \sim N(0,1)$ are i.i.d. Gaussians so that $\mathbb{E}[ZZ^\top] = \sum_{j=1}^n v_j v_j^\top$ has trace $\mathrm{Tr}\big[\mathbb{E}[ZZ^\top]\big] = \sum_{j=1}^n \|v_j\|_2^2 \leq n$. Let $u_1, \ldots, u_r$ be an orthonormal basis of H

with $r \leq n$, and write $\sum_{j=1}^n v_j v_j^\top = \sum_{j=1}^r \sigma_j u_j u_j^\top$. Since $\sum_{j=1}^n v_j v_j^\top \succeq 0$, we have $\sigma_j \geq 0$ for all j. Then after reindexing we may assume that $0 \leq \sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_r$. Since $\sum_{j=1}^r \sigma_j = \sum_{j=1}^n \|v_j\|_2^2 \leq n$ we know by Markov's Inequality that $\sigma_{2n/3} \leq 3/2$, denoting $\sigma_j = 0$ for $j > r$. Thus restricting to the subspaces $F := \mathrm{span}\{u_1, \ldots, u_{2n/3}\}$ and $V := \{g \in \mathbb{R}^n \mid \sum_{j=1}^n g_j v_j \in F\}$ with $\dim(V) \geq \frac{2}{3}n$, we may lower bound

$$
\begin{aligned}
\Pr_{g \sim N(0, I_V)}\Big[\sum_{j=1}^n g_j v_j \in 3/2 \cdot K\Big] &= \Pr_{g \sim N(0, I_{2n/3})}\Big[\sum_{j=1}^{2n/3} g_j \cdot \sigma_j u_j u_j^\top \in 3/2 \cdot K\Big] \\
&\overset{(*)}{\geq} \Pr_{g \sim N(0, I_{2n/3})}\Big[\sum_{j=1}^{2n/3} g_j \cdot 3/2 \cdot u_j u_j^\top \in 3/2 \cdot K\Big] \\
&= \gamma_F(K \cap F) \\
&\overset{\text{Lem 20}}{\geq} \gamma_H(K \cap H) \\
&\geq e^{-\alpha n},
\end{aligned}
$$

where $(*)$ follows by Lemma 21. Then by Theorem 78, the symmetric convex body $Q := \{x \in \mathbb{R}^n : \sum_{j=1}^n x_j v_j \in K\}$ contains a good partial coloring in $Q \cap F$. $\square$

Then Lemma 87 implies the existence of a partial coloring with optimal bounds as long as n is of the order of d:

Corollary 88. *Let $K \subseteq \mathbb{R}^d$ be a normalized zonotope and let $v_1, \ldots, v_n \in K$. Then there is a randomized polynomial time algorithm to find a good partial coloring $x \in [-1, 1]^n$ so that $\|\sum_{j=1}^n x_j v_j\|_K \lesssim \sqrt{d}$.*

Proof. By Theorem 75, denoting $H := \mathrm{span}\{v_1, \ldots, v_n\}$, we have $\gamma_H(C \cdot K \cap H) \geq e^{-n}$. By Lemma 16, there exists some constant $\alpha > 0$ such that $\gamma_H(K \cap H) \geq e^{-\alpha n}$. By Lemma 82, $v_i \in \sqrt{d} B_2^d$, so that the statement follows directly from Lemma 87. $\square$

3.4.2 Proof of the main Theorem

Now we have all the ingredients to prove our main result, Theorem 73.

Proof of Theorem 73. By Theorem 80, we may assume that K is generated by only $m \lesssim d \log d$ segments, and by Lemma 81, we may assume that K is a normalized zonotope $K := \sqrt{\frac{d}{m}} A^\top B_\infty^m$ for some approximately regular $A \in \mathbb{R}^{m \times d}$. By Theorem 3, since $\mathsf{vb}(K,K) \leq 2 \cdot \mathsf{vb}_d(K,K)$, we may assume that $n = d$, though for clarity we only use this in the final bound. As before we set $Q := \{x \in \mathbb{R}^n : \sum_{j=1}^n x_j v_j \in K\}$. We iteratively apply Lemma 87 for t rounds to obtain a partial coloring $x' \in Q \cap [-1,1]^n$, so that the set $I := \{i : |x'_i| < 1\}$ of partially colored indices satisfies $|I| \leq n/2^t$, and by the triangle inequality over the t rounds $\| \sum_{j=1}^n x'_j v_j \|_K \lesssim \sqrt{d} \cdot t$.

For each $j \in I$, we may write $v_j = \sqrt{\frac{d}{m}} A^\top u_i$ for some $u_i \in B_\infty^m$. By Theorem 77, we can find $\tilde{x} \in \mathbb{R}^n$ so that $x := \tilde{x} + x' \in \{-1,1\}^n$ and $\sum_{i \in I} \tilde{x}_i u_i \in \sqrt{|I| \log(\frac{2m}{|I|})} \cdot c \cdot B_\infty^m$ where we set $\tilde{x}_i = 0$ for $i \notin I$. Therefore, setting $t := \log\log(\frac{2m}{n})$,

$$
\begin{aligned}
\Big\| \sum_{j=1}^n x_j v_j \Big\|_K &\leq \Big\| \sum_{j=1}^n x'_j v_j \Big\|_K + \Big\| \sum_{j \in I} \tilde{x}_j v_j \Big\|_K \\
&\lesssim \sqrt{d} \cdot t + \sqrt{\frac{n}{2^t} \cdot \log\Big(\frac{2m}{n/2^t}\Big)} \\
&= \sqrt{d} \log\log\Big(\frac{2m}{n}\Big) + \underbrace{\sqrt{\frac{n}{\log(\frac{2m}{n})} \cdot \log\Big(\frac{2m}{n} \cdot \log\Big(\frac{2m}{n}\Big)\Big)}}_{\lesssim \sqrt{n} \leq \sqrt{d}} \\
&\lesssim \sqrt{d} \log\log\Big(\frac{2m}{n}\Big) \\
&\lesssim \sqrt{d} \log\log\Big(\frac{d \log d}{n}\Big).
\end{aligned}
$$

We conclude that $\mathsf{vb}(K,K) \lesssim \mathsf{vb}_d(K,K) \lesssim \sqrt{d} \log\log\log d$. $\square$

3.5 The vector balancing constant $\mathrm{vb}(K,Q)$

In this section we prove Theorem 76, stating that $\mathsf{vb}(K,Q) \lesssim \sqrt{d \log d}$ where K and Q are normalized zonotopes. First note that Cor 88 indeed generalizes and for any $v_1, \ldots, v_n \in K$ there is a good partial coloring $x \in [-1,1]^n$ with $\| \sum_{j=1}^n x_j v_j \|_Q \lesssim \sqrt{d}$. On the other hand, in the proof of Theorem 73 we have also relied on Spencer's Theorem which implies that

$\mathsf{vb}_n(K,K) \lesssim \sqrt{n\log(\frac{2m}{n})}$. In particular this gives a bound that improves as n decreases. However in our setting with different zonotopes K and Q such a bound does not hold!

To see this, let $H \in \{-1,1\}^{d\times d}$ be a Hadamard matrix, meaning that all rows and columns are orthogonal. Then one can verify that $K := \frac{1}{\sqrt{d}}H^\top B_\infty^d$ is a normalized zonotope; in fact, K is a rotated cube. Fix any $n \leq d$ and consider the points $v_1,\ldots,v_n \in K$ with $v_i = \frac{1}{\sqrt{d}}H^\top H^i = \sqrt{d}\cdot e_i$. We choose $Q := B_\infty^d$ as the second normalized zonotope. Any good partial coloring $x \in [-1,1]^n$ must have a coordinate i with $|x_i| \geq \frac{1}{2}$ and so $\|\sum_{j=1}^n x_j v_j\|_Q \geq \sqrt{d}|x_i| \geq \frac{\sqrt{d}}{2}$.

Hence instead of applying Cor 88 iteratively and obtaining a bound of $\mathsf{vb}(K,Q) \lesssim \sqrt{d}\log d$, we use Banaszczyk's Theorem together with Theorem 75:

Proof of Theorem 76. Let $K,Q \subseteq \mathbb{R}^d$ be normalized zonotopes, and let $v_1,\ldots,v_n \in K$ be the vectors to be balanced. Define $H := \mathrm{span}\{v_1,\ldots,v_n\}$ and let $r := \dim(H) \leq \min\{d,n\}$. By applying Theorem 75 to the zonotope Q, subspace H, and $t := \sqrt{2\log 2r}$, we find that

$$\gamma_H\big(\sqrt{2\log 2r}C'Q \cap H\big) \geq e^{-\frac{1}{2}} > \frac{1}{2}.$$

By Lemma 82 we know that $v_i \in \sqrt{d}B_2^d$ for each $i \in [n]$, hence by Theorem 12, signs $x \in \{-1,1\}^n$ can be computed in polynomial time such that

$$\sum_{j=1}^n x_j v_j \in \sqrt{d}C''\left(\sqrt{2\log 2r}C'Q \cap H\right) \subseteq C\sqrt{d\log\min\{d,n\}}Q,$$

as desired. In particular, $\mathrm{vb}(K,Q) \lesssim \sqrt{d\log d}$. $\square$

3.6 Open problems

The main open question about zonotopes is whether a d-dimensional zonotope can be approximated up to a constant factor using only a linear number of segments:

Conjecture 9 ([152])**.** *For any zonotope $K \subseteq \mathbb{R}^d$ and $0 < \varepsilon \leq \frac{1}{2}$, does there exist a zonotope Q with $O(\frac{d}{\varepsilon^2})$ segments so that $Q \subseteq K \subseteq (1+\varepsilon)Q$?*

Equivalently, since the polar body of a zonotope $A^\top B_\infty^m \subseteq \mathbb{R}^d$ is the preimage $A^{-1}(B_1^m) := \{x \in \mathbb{R}^d : \|Ax\|_1 \leq 1\}$, we can restate the question as follows:

Conjecture 10. *Does there exist a universal constant $C > 0$ such that given any matrix $A \in \mathbb{R}^{m\times d}$ with $m \geq d$ and $0 < \varepsilon \leq \frac{1}{2}$, one can always find another matrix $\tilde{A} \in \mathbb{R}^{Cd/\varepsilon^2 \times d}$ with $\|\tilde{A}x\|_1 \leq \|Ax\|_1 \leq (1+\varepsilon)\|\tilde{A}x\|_1$ for all $x \in \mathbb{R}^d$?*

We remark that if one replaces the ℓ_1 norm by the ℓ_2 norm, an analogue of Conjecture 10 holds as a direct corollary of a linear-size spectral sparsifier [31]. In that setting, each row of $\tilde{A}$ is a scalar multiple of a row of A, and there is hope that another rescaling of the rows of A may suffice for the ℓ_1 norm. Just as a spectral sparsifier can be found via spectral partial colorings [138], we also state the stronger conjecture of the existence of good partial colorings in the ℓ_1 setting:

Conjecture 11. *Given any matrix $A \in \mathbb{R}^{m\times d}$, does the set*

$$K := \Big\{x \in \mathbb{R}^m : \Big|\sum_{i=1}^m x_i |\langle A_i, z\rangle|\Big| \leq \sqrt{\tfrac{d}{m}}\|Az\|_1 \ \forall z \in \mathbb{R}^d\Big\}$$

have large Gaussian measure $\gamma_m(K) \geq e^{-Cm}$ where $C > 0$ is a universal constant?

Finally, we restate Schechtman's question, which would also follow from the above conjectures:

Conjecture 12 ([152])**.** *Is it true that for any zonotope $K \subseteq \mathbb{R}^d$, $\mathsf{vb}(K,K) \lesssim \sqrt{d}$?*

3.7 Normalizing zonotopes

In this section, we show that for any full-dimensional zonotope $K \subseteq \mathbb{R}^d$ there is a linear transformation $T : \mathbb{R}^d \to \mathbb{R}^d$ and a normalized zonotope $\tilde{K}$ so that $\frac{4}{5}\tilde{K} \subseteq T(K) \subseteq \tilde{K}$. For this result we will need the existence of *Lewis weights* [47]:

Theorem 89. *Given a matrix $A \in \mathbb{R}^{m\times d}$, there exists a unique vector $\overline{w} \in \mathbb{R}^m_{>0}$ so that for all $i \in [m]$ one has*

$$\overline{w}_i^{-2} A_i^\top (A^\top \overline{W}^{-1} A)^{-1} A_i = 1,$$

where $\overline{W} := \mathrm{diag}(\overline{w})$*. Moreover,* $\mathrm{Tr}[\overline{W}] \leq d$*, with equality for full rank* A*.*

Now to the proof of Lemma 81.

Proof of Lemma 81. Consider a full-dimensional zonotope $K = A^\top B_\infty^m$ with $A \in \mathbb{R}^{m\times d}$. Let $\overline{W}$ be the diagonal matrix corresponding to the Lewis weights of A and let $W := D\overline{W}$ where $D > 0$ is large enough so that $w_i := W_{i,i} \geq 1$ for all i. Define a matrix $B := A(A^\top W^{-1}A)^{-1/2} \in \mathbb{R}^{m\times d}$ and define a second matrix $\tilde{A}$ where each row B_i is replaced by $\lceil w_i \rceil$ many rows so that the first $\lfloor w_i \rfloor$ rows are all copies of $w_i^{-1}B_i$, and (if $\{w_i\} \neq 0$) the last row is $\{w_i\}^{1/2} w_i^{-1} B_i$, for a total of $m' := \sum_{i=1}^m \lceil w_i \rceil$ many rows. We will show that the conditions of Lemma 81 hold with

$$T : \mathbb{R}^d \to \mathbb{R}^d,\ \ T(K) = \sqrt{\tfrac{d}{m'}}(A^T W^{-1} A)^{-1/2} K = \sqrt{\tfrac{d}{m'}} \cdot B^T B_\infty^m$$

and $\widetilde{K} := \sqrt{\frac{d}{m'}} \widetilde{A}^T B_\infty^{m'}$.

First we show that $\tilde{K}$ is normalized, or equivalently that $\tilde{A}$ is approximately regular. Note that

$$(\tilde{A}^\top \tilde{A})_{j,k} = \sum_{i=1}^{m'} \tilde{A}_{i,j}\tilde{A}_{i,k} = \sum_{i=1}^{m} w_i^{-2}(\lfloor w_i \rfloor + \{w_i\}) \cdot B_{i,j}B_{i,k} = \sum_{i=1}^{m} w_i^{-1} \cdot B_{i,j}B_{i,k},$$

so that by definition of B,

$$\tilde{A}^\top \tilde{A} = B^\top W^{-1} B = (A^\top W^{-1} A)^{-1/2} A^\top W^{-1} A (A^\top W^{-1} A)^{-1/2} = I_d.$$

Moreover, by the definition of Lewis weights, for each row $i' \in [m']$ corresponding to a copy of B_i one has

$$\|\tilde{A}_{i'}\|_2^2 \leq w_i^{-2} B_i^\top B_i = w_i^{-2} A_i^\top (A^\top W^{-1} A)^{-1} A_i = \frac{1}{D} \leq \frac{2d}{m'},$$

where the last inequality follows since

$$m' = \sum_{i=1}^m \lceil w_i \rceil \leq 2 \cdot \sum_{i=1}^m w_i = 2D \sum_{i=1}^m \overline{w}_i \leq 2D \cdot d.$$

Thus $\widetilde{A}$ is approximately regular, and $\tilde{K}$ is normalized.

To see that $\frac{4}{5}\tilde{K} \subseteq T(K) \subseteq \tilde{K}$, take an arbitrary

$$y = \tfrac{4}{5}\sqrt{\tfrac{d}{m'}}\sum_{i=1}^{m}\Big(\sum_{p=1}^{\lfloor w_i \rfloor} x_{i,p} w_i^{-1} B_i + x_{i,\lceil w_i \rceil}\{w_i\}^{1/2} w_i^{-1} B_i\Big) \in \tfrac{4}{5}\sqrt{\tfrac{d}{m'}}\tilde{A}^{\top} B_{\infty}^{m'} = \tfrac{4}{5}\tilde{K},$$

and rewrite it as

$$\tfrac{4}{5}\sqrt{\tfrac{d}{m'}}\sum_{i=1}^{m}\Big(\underbrace{w_i^{-1}\Big(\sum_{i=1}^{\lfloor w_i \rfloor} x_{i,p} + x_{i,\lceil w_i \rceil}\{w_i\}\Big)}_{\in[-1,1]} + \underbrace{x_{i,\lceil w_i \rceil}\tfrac{\{w_i\}^{1/2} - \{w_i\}}{w_i}}_{\in[-\frac{1}{4},\frac{1}{4}]}\Big) B_i \in \sqrt{\tfrac{d}{m'}} B^{\top} B_{\infty}^{m} = T(K).$$

Now taking an arbitrary $y := \sqrt{\frac{d}{m'}}\sum_{i=1}^{m} x_i B_i \in B^{\top} B_{\infty}^{m} = T(K)$, we may write

$$y = \sqrt{\tfrac{d}{m'}}\sum_{i=1}^{m}\Big(\sum_{p=1}^{\lfloor w_i \rfloor} x_i w_i^{-1} B_i + x_i\{w_i\} w_i^{-1} B_i\Big) \in \sqrt{\tfrac{d}{m'}}\tilde{A}^{\top} B_{\infty}^{m'} = \tilde{K},$$

completing the proof of the lemma. Finally, note that this result immediately implies that

$$\tfrac{4}{5}\mathsf{vb}(\tilde{K},\tilde{K}) \leq \mathsf{vb}(K,K) \leq \tfrac{5}{4}\mathsf{vb}(\tilde{K},\tilde{K}). \qquad \square$$

Chapter 4

APPROXIMATE CARATHÉODORY AND VECTOR BALANCING

This chapter is based on a joint paper with Thomas Rothvoss [139].

4.1 Introduction

The (exact) Carathéodory Theorem is part of most introductory courses on the theory of linear programming: given any vector $\boldsymbol{z} \in \text{conv}(X)$ where $X \subseteq \mathbb{R}^m$, there is a subset of points $X' \subseteq X$ with $|X'| \leq m+1$ so that already $\boldsymbol{z} \in \text{conv}(X')$. More recently, the *approximate* version gained interest, where only k vectors from X may be selected with uniform weights and the goal is to minimize the error in a given norm.

Barman [29] used an approximate Carathéodory bound for algorithms to compute approximate Nash equilibria for bimatrix games as well as for finding k-densest subgraphs. The core argument of [29] is as follows: if one has two players with n strategies each and some payoff matrix $\boldsymbol{A} \in [-1,1]^{n \times n}$, then for any mixed strategy $\boldsymbol{y}$ of the column player (i.e. $\boldsymbol{y} \in \mathbb{R}^n_{\geq 0}$ and $\|\boldsymbol{y}\|_1 = 1$) one can apply the approximate Carathéodory Theorem for norm $\|\cdot\|_\infty$ (or rather equivalently for $\|\cdot\|_{\log(n)}$) and find $k := \Theta(\frac{\log(n)}{\varepsilon^2})$ columns $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k$ of $\boldsymbol{A}$ so that $\|\boldsymbol{A}\boldsymbol{y} - \frac{1}{k}\sum_{i=1}^k \boldsymbol{a}_i\|_\infty \leq \varepsilon$. In other words, any mixed strategy can be ε-approximated by the unweighted average of only $\Theta(\frac{\log(n)}{\varepsilon^2})$ many pure strategies which then allows for an efficient enumeration. The approximate Carathéodory Theorem has also been useful in algebraic settings. For example Deligkas et al [59] use it to find approximate solutions to systems of polynomial (in)equalities and Bhargava, Saraf and Volkovich [34] use approximate Carathéodory to prove that sparse polynomials have only sparse factors which then allows efficient deterministic factorization of sparse polynomials; both applications use the variant with respect to the $\|\cdot\|_\infty$-norm.

To make the statements formal, for symmetric convex bodies $P, Q \subseteq \mathbb{R}^m$ and $k \in \mathbb{N}$, we denote

$$\mathrm{ac}_k(P, Q) := \sup_{\substack{X \subseteq P,\\ \boldsymbol{z} \in \mathrm{conv}(X)}} \inf_{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in X} \Big\| \boldsymbol{z} - \frac{1}{k} \sum_{i=1}^{k} \boldsymbol{v}_i \Big\|_Q$$

as the best error bound with respect to the $\| \cdot \|_Q$-norm for approximating a point $\boldsymbol{z}$ in the convex hull of some points in P. We would like to point out that the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ may be taken with repetition. Here, $\| \cdot \|_Q$ is the norm with $\|\boldsymbol{x}\|_Q = \min\{s \geq 0 \mid \boldsymbol{x} \in sQ\}$. A folklore result is that for the Euclidean norm one has $\mathrm{ac}_k(B_2^m, B_2^m) \leq \frac{1}{\sqrt{k}}$ for any $k \geq 1$, which gives a *dimension free* bound. More generally, for $p \geq 1$, it is true that $\mathrm{ac}_k(B_p^m, B_p^m) \leq O(\sqrt{\frac{p}{k}})$ where $B_p^m := \{\boldsymbol{x} \in \mathbb{R}^m \mid \|\boldsymbol{x}\|_p \leq 1\}$ is the $\| \cdot \|_p$-unit ball. This bound is derived from Maurey's Lemma from functional analysis (which was reported by Pisier [132]; for an English version, see the appendix of Bourgain and Nelson [37]). Algorithmically, the result is simple: write $\boldsymbol{z} = \sum_{i=1}^{N} \lambda_i \boldsymbol{u}_i$ where $\lambda_1, \ldots, \lambda_N \geq 0$ and $\sum_{i=1}^{N} \lambda_i = 1$. Then sample $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_N\}$ independently according to the probabilities λ_i (possibly with repetition).

Another approach in the literature by Mirrokni et al [118] is based on the desire to avoid the computation of $\boldsymbol{z} = \sum_{i=1}^{N} \lambda_i \boldsymbol{u}_i$. Instead they use the *Mirror Descent* algorithm from convex optimization to compute the sequence $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ directly. In fact they reprove the bound of $\mathrm{ac}_k(B_p^m, B_p^m) \leq O(\sqrt{\frac{p}{k}})$ using their framework. More recently, Combettes and Pokutta [48] show that the Frank-Wolfe algorithm can also be used to recover the same bounds. From the current state of the literature, there are two directions that appear natural to follow:

- *Approximate Carathéodory for general pairs of norms.* The existing bounds are for the case where P and Q are the same $\| \cdot \|_p$-ball. Is there a convenient framework that can handle general symmetric convex bodies or at least $P = B_p^m$ and $Q = B_q^m$?

- *Tight bounds for approximate Carathéodory.* Generally, it is stated that for example the bound $\mathrm{ac}_k(B_p^m, B_p^m) \leq O(\sqrt{\frac{p}{k}})$ is tight (see e.g. [118]). But that is only true if one aims for a dimension independent bound. So for which regimes of m vs. k and p is it possible to improve the bound?

A classical area within combinatorics that appears related to these questions is *discrepancy theory*. Let $S_1, \dots, S_m \subseteq \{1, \dots, n\}$ be a set system over n elements. Then the goal is to find a bi-coloring $\boldsymbol{x} \in \{-1,1\}^n$ so that the worst imbalance $\max_{i=1,\dots,m} |\sum_{j \in S_i} x_j|$ is minimized. A seminal result of Spencer [157] says that for $m \geq n$, the discrepancy is bounded by $O(\sqrt{n \log(\frac{2m}{n})})$. If no element is in more than t sets, then one can also prove a bound of $2t$, see Beck and Fiala [33]. A convex geometry based method by Banaszczyk [13] shows that for any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ with column length $\|\boldsymbol{A}^j\|_2 \leq 1$ for all $j = 1, \dots, n$ and any symmetric convex body $K \subseteq \mathbb{R}^m$ with Gaussian measure at least $1/2$ (for example $K = \Theta(\sqrt{\log(m)}) \cdot B_\infty^m$ or $K = \Theta(\sqrt{m}) \cdot B_2^m$ work), there is a coloring $\boldsymbol{x} \in \{-1,1\}^n$ with $\boldsymbol{Ax} \in 5K$. Interestingly, neither of these cited results of [157, 33, 13] can be obtained by merely taking a uniform random coloring $\boldsymbol{x}$. But for example, the result by Spencer allows for elegant algorithmic proofs [16] [103].

4.1.1 *Our contribution*

The connection between the approximate Carathéodory problem and vector balancing was already discovered by Dadush et al [53] who proved that for any symmetric convex bodies $P, Q \subseteq \mathbb{R}^m$ one has $\mathrm{ac}_k(P,Q) \leq \frac{\mathrm{vb}(P,Q)}{k}$. But for example for $P = Q = B_2^m$ one has $\mathrm{vb}(B_2^m, B_2^m) = \Theta(\sqrt{m})$ and so the obtained bound is $O(\frac{\sqrt{m}}{k})$, which is suboptimal if $k \ll m$. Instead, we prove the following:

Theorem 90. *For any symmetric convex bodies $P, Q \subseteq \mathbb{R}^m$ and any $k \in \mathbb{N}$ one has*

$$ac_k(P,Q) \leq 4 \sum_{\ell \geq 1} \frac{1}{k 2^\ell} \cdot \mathsf{vb}_{k2^\ell}(P,Q)$$

The vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ can be found in time $O(\log m)$ times the time to find a coloring behind $vb_t(P,Q)$ where $t \leq m+1$, assuming we are given $\boldsymbol{z}$ as convex combination of at most $m+1$ vectors from X.

For most bodies the quantity $\mathrm{vb}_k(P,Q)$ grows sublinear in k and the sum is dominated by the first term, in which case one has $\mathrm{ac}_k(P,Q) \lesssim \frac{1}{k} \cdot \mathrm{vb}_k(P,Q)$. For example if $P = Q = B_2^m$ one then has $\mathrm{vb}_k(B_2^m, B_2^m) = \Theta(\sqrt{k})$ and so one recovers the $O(\frac{1}{\sqrt{k}})$ bound mentioned earlier[1]. Also note that by [102], for any symmetric convex bodies $P, Q \subseteq \mathbb{R}^m$ and any $t \in \mathbb{N}$ one has $\mathrm{vb}_t(P,Q) \leq 2\mathrm{vb}_m(P,Q)$, meaning that the worst case is basically attained for m many vectors. Then the infinite sum in Theorem 90 is dominated by the first $\log(2m/k)$ terms if $k \leq m$ and the first term if $k \geq m$.

Combining Theorem 90 above with Theorem 52 then gives:

Theorem 91. *Let $2 \leq p \leq q \leq \infty$ and $k \in \mathbb{N}$. Then*

$$ac_k(B_p^m, B_q^m) \lesssim \frac{1}{\frac{1}{2} - \frac{1}{p} + \frac{1}{q}} \cdot \frac{\sqrt{\min\{p, \log(\frac{2m}{k})\}}}{k^{1/2+1/p-1/q}}$$

The vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ can be found in time $O(\log^2 m) \cdot T(m, m+1) \leq O(m^{1+\omega} \log^2(m))$ assuming we are given $\boldsymbol{z}$ as a convex combination of at most $m+1$ vectors in X.

To the best of our knowledge this is the first approximate Carathéodory bound for pairs of different $\|\cdot\|_p$-norms. In particular, when $p = q$ this improves upon the $O(\sqrt{\frac{p}{k}})$ bound in [118] whenever $p \ll \log(\frac{2m}{k})$:

Corollary 92. *Let $2 \leq p \leq \infty$ and $k \in \mathbb{N}$. Then*

$$ac_k(B_p^m, B_p^m) \lesssim \sqrt{\frac{\min\{p, \log(\frac{2m}{k})\}}{k}}$$

The vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ can be found in polynomial time assuming we are given $\boldsymbol{z}$ as a convex combination of at most $m+1$ vectors in X.

[1]Though not all bodies allow for such a sublinear dependence. For example fix $1 \leq k \leq m/2$. Then one has $\mathrm{vb}_k(B_1^m, B_1^m) = \Theta(k)$ and so Theorem 90 provides a suboptimal bound of $\mathrm{ac}_k(B_1^m, B_1^m) \leq O(\log \frac{m}{k})$. On the other hand, indeed it is true that $\mathrm{ac}_k(B_1^m, B_1^m) \geq \Omega(1)$ meaning that the approximate Carathéodory bound does not even improve with k (at least as long as $k \leq m/2$). To see this, consider $X := \{\boldsymbol{e}_1, \dots, \boldsymbol{e}_m\}$ and a target of $\boldsymbol{z} := (\frac{1}{m}, \dots, \frac{1}{m})$. Then for any $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k \in X$ one has $\|\boldsymbol{z} - \frac{1}{k}\sum_{i=1}^k \boldsymbol{v}_i\|_1 \geq \Omega(1)$ since any coordinate whose unit vector is not included will contribute $\frac{1}{m}$ to the norm.

4.2 Preliminaries

In this section we review a few facts that we later rely on. Let $S^{m-1} := \{\boldsymbol{x} \in \mathbb{R}^m \mid \|\boldsymbol{x}\|_2 = 1\}$ be the sphere.

Convex functions. Recall the following well known fact:

Lemma 93 (Jensen Inequality for convave functions)**.** *Let X be any $\mathbb{R}$-valued random variable and let $F : \mathbb{R} \to \mathbb{R}$ be a concave function, then $F(\mathbb{E}[X]) \geq \mathbb{E}[F(X)]$.*

Estimates on $\|\cdot\|_p$ norms. It will be useful to understand how the norm $\|\boldsymbol{z}\|_p$ of a vector can change depending on $p \in [1, \infty]$.

Lemma 94. *For any $\boldsymbol{z} \in \mathbb{R}^m$ and $1 \leq p \leq q \leq \infty$ one has $\|\boldsymbol{z}\|_q \leq \|\boldsymbol{z}\|_p \leq m^{1/p-1/q}\|\boldsymbol{z}\|_q$.*

Lemma 95. *For any $\boldsymbol{z} \in \mathbb{R}^m$ and $1 \leq p \leq q \leq \infty$, we have $\|\boldsymbol{z}\|_q^q \leq \|\boldsymbol{z}\|_p^p \cdot \|\boldsymbol{z}\|_\infty^{q-p}$.*

4.3 Reduction from Approximate Carathéodory to Vector Balancing

In this section, we prove the reduction of the approximate Carathéodory problem to vector balancing as stated in Theorem 90. The idea is to follow the classical approach of [102]: begin with an arbitrary convex combination and round the coefficients bit-by-bit. The same basic approach was also followed by Dadush et al [53]. We prove an auxiliary lemma that bounds the error when "doubling the fractionality".

Lemma 96. *Let $P, Q \subseteq \mathbb{R}^m$ be symmetric convex bodies and let $\delta > 0$. Let $\boldsymbol{z} = \sum_{i=1}^n \lambda_i \boldsymbol{v}_i$ where $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in P$ and $\boldsymbol{\lambda} \in \delta\mathbb{Z}_{\geq 0}^n$. Then there is a vector $\boldsymbol{z}' = \sum_{i=1}^n \lambda_i' \boldsymbol{v}_i$ where $\boldsymbol{\lambda}' \in 2\delta\mathbb{Z}_{\geq 0}^n$ so that $\|\boldsymbol{z} - \boldsymbol{z}'\|_Q \leq \delta \cdot vb_n(P, Q)$ and $\sum_{i=1}^n \lambda_i' \leq \sum_{i=1}^n \lambda_i$.*

Proof. Write $\lambda_i = 2\delta a_i + \delta b_i$ with $b_i \in \{0, 1\}$ and $a_i \in \mathbb{Z}_{\geq 0}$. Let $I := \{i \in [n] \mid b_i = 1\}$. Now, let $\boldsymbol{x} \in \{-1, 1\}^I$ be the coloring so that $\|\sum_{i \in I} x_i \boldsymbol{v}_i\|_Q \leq \mathrm{vb}_{|I|}(P, Q) \leq \mathrm{vb}_n(P, Q)$. We may assume that $\sum_{i \in I} x_i \leq 0$ — otherwise replace $\boldsymbol{x}$ with $-\boldsymbol{x}$. We may extend the vector to

$\boldsymbol{x} \in \{-1, 0, 1\}^m$ by setting $x_i := 0$ for $i \notin I$. We update $\lambda'_i := 2\delta a_i + \delta(1 + x_i) b_i \in 2\delta\mathbb{Z}_{\geq 0}$ for $i \in [n]$. Next, we define $\boldsymbol{z}' := \sum_{i=1}^n \lambda'_i \boldsymbol{v}_i$. Then

$$\|\boldsymbol{z} - \boldsymbol{z}'\|_Q = \delta \Big\| \sum_{i \in I} x_i \boldsymbol{v}_i \Big\|_Q \leq \delta \cdot \mathrm{vb}_n(P, Q)$$

Note that $\sum_{i \in I} x_i \leq 0$ implies that $\sum_{i=1}^n \lambda'_i \leq \sum_{i=1}^n \lambda_i$. This gives the claim. □

Next, we iteratively apply Lemma 96 to an initial convex combination until the convex coefficients are multiples of $\frac{1}{k}$. We almost obtain the desired claim, just that the number of vectors might be *less* than k.

Lemma 97. *Let $P, Q \subseteq \mathbb{R}^m$ be symmetric convex bodies. Then for any $\boldsymbol{z} \in conv(X)$ with $X \subseteq P$ and $k \in \mathbb{N}$ there are $s \in \{0, \ldots, k\}$ and $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_s \in X$ so that*

$$\Big\| \boldsymbol{z} - \frac{1}{k} \sum_{i=1}^s \boldsymbol{v}_i \Big\|_Q \leq \sum_{\ell \geq 1} \frac{2}{k 2^\ell} \cdot vb_{k2^\ell}(P, Q)$$

The vectors can be found in time $O(\log m)$ times the time to find the colorings in $vb_t(P, Q)$ where $t \leq m + 1$.

Proof. Fix a point $\boldsymbol{z} \in \mathrm{conv}(X)$ where $X \subseteq P$. Then we can write $\boldsymbol{z} = \sum_{i=1}^n \lambda_i \boldsymbol{v}_i$ where $n \leq m + 1$, $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in X$, $\lambda_i \geq 0$ for all $i = 1, \ldots, m$ and $\sum_{i=1}^m \lambda_i = 1$. Without loss of generality we may assume that $\boldsymbol{\lambda} \in \frac{2^{-L}}{k} \mathbb{Z}_{\geq 0}^n$ for some $L \in \mathbb{N}$. We abbreviate $\boldsymbol{z}^{(L)} := \boldsymbol{z}$. Now suppose for $\ell \in \{0, \ldots, L\}$ the current iterate is $\boldsymbol{z}^{(\ell)}$ so that $\boldsymbol{z}^{(\ell)} = \sum_{i=1}^n \lambda_i^{(\ell)} \boldsymbol{v}_i$ and $\boldsymbol{\lambda}^{(\ell)} \in \frac{2^{-\ell}}{k} \mathbb{Z}_{\geq 0}^n$.

Then we apply Lemma 96 to obtain a vector $\boldsymbol{z}^{(\ell-1)} = \sum_{i=1}^n \lambda_i^{(\ell-1)} \boldsymbol{v}_i$ with $\boldsymbol{\lambda}^{(\ell-1)} \in \frac{2^{-(\ell-1)}}{k} \mathbb{Z}_{\geq 0}^n$ and $\sum_{i=1}^n \lambda_i^{(\ell-1)} \leq 1$. Using that $|\mathrm{supp}(\boldsymbol{\lambda}^{(\ell)})| \leq k 2^\ell$, the approximation error satisfies

$$\big\| \boldsymbol{z}^{(\ell)} - \boldsymbol{z}^{(\ell-1)} \big\|_Q \leq \frac{1}{2^\ell k} \cdot \mathrm{vb}_{|\mathrm{supp}(\boldsymbol{\lambda}^{(\ell)})|}(P, Q) \leq \frac{1}{2^\ell k} \cdot \mathrm{vb}_{k 2^\ell}(P, Q)$$

Note that the final iterate is of the form $\boldsymbol{z}^{(0)} = \sum_{i=1}^n \lambda_i^{(0)} \boldsymbol{v}_i$ with $\lambda_i^{(0)} \in \frac{\mathbb{Z}_{\geq 0}}{k}$ and $\sum_{i=1}^n \lambda_i^{(0)} \leq 1$. Then for $s := k \sum_{i=1}^n \lambda_i^{(0)} \in \{0, \ldots, k\}$, let $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_s$ be a list of vectors that contains $\boldsymbol{v}_i$ exactly $k \lambda_i^{(0)} \in \mathbb{Z}_{\geq 0}$ times. Using the triangle inequality we obtain

$$\Big\| \boldsymbol{z} - \frac{1}{k} \sum_{i=1}^s \boldsymbol{u}_i \Big\|_Q \leq \sum_{\ell=1}^L \|\boldsymbol{z}^{(\ell)} - \boldsymbol{z}^{(\ell-1)}\|_Q \leq \sum_{\ell \geq 1} \frac{1}{k 2^\ell} \cdot \mathrm{vb}_{k 2^\ell}(P, Q)$$

Now let us discuss the running time. First, we can choose $L \leq O(\log m)$ while possibly making a rounding error of $\max\{\|\boldsymbol{y}\|_Q : \boldsymbol{y} \in P\} \leq \mathrm{vb}_1(P,Q)$, which we absorb by paying an extra factor of 2. Then the running time is dominated by the time to find the colorings. Note that we call $\mathrm{vb}_t(P,Q)$ only L times for parameters t with $t \leq m+1$. $\square$

Now we will use the same trick as Dadush et al [53] in order to obtain exactly k vectors, at the expense of a factor 2 in the approximation error:

Proof of Theorem 90. Let $\boldsymbol{z} \in \mathrm{conv}(X)$ where $X \subseteq P$. Fix a vector $\boldsymbol{u}_0 \in X$ and write $X' := \{\boldsymbol{u} - \boldsymbol{u}_0 \mid \boldsymbol{u} \in X\}$. Note that in particular $\mathbf{0} \in X'$ and $\boldsymbol{z} - \boldsymbol{u}_0 \in \mathrm{conv}(X')$. We apply Lemma 97 and obtain vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_s \in X'$ with $s \leq k$ so that

$$\Big\|(\boldsymbol{z} - \boldsymbol{u}_0) - \frac{1}{k}\sum_{i=1}^{s} \boldsymbol{v}_i\Big\|_Q \leq \sum_{\ell \geq 1} \frac{2}{k2^\ell} \cdot \mathrm{vb}_{k2^\ell}(2P, Q)$$

using that $X' \subseteq 2P$. Since $\mathbf{0} \in X'$, we can extend this sequence to a list $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ of k vectors. Each $\boldsymbol{v}_i \in X'$ can be written as $\boldsymbol{v}_i = \boldsymbol{u}_i - \boldsymbol{u}_0$ with $\boldsymbol{u}_1, \dots, \boldsymbol{u}_k \in X$. Then

$$\Big\|\boldsymbol{z} - \frac{1}{k}\sum_{i=1}^{k} \boldsymbol{u}_k\Big\|_Q = \Big\|(\boldsymbol{z} - \boldsymbol{u}_0) - \frac{1}{k}\sum_{i=1}^{k} (\boldsymbol{u}_k - \boldsymbol{u}_0)\Big\|_Q \leq 4\sum_{\ell \geq 1} \frac{1}{k2^\ell} \cdot \mathrm{vb}_{k2^\ell}(P, Q). \qquad \square$$

4.4 *Approximate Carathéodory bounds for ℓ_p norms*

Next, we prove the bound on $\mathrm{ac}_k(B_p^m, B_q^m)$ claimed in Theorem 91:

Proof of Theorem 91. Let $2 \leq p \leq q \leq \infty$. We apply the reduction to the vector balancing constant from Theorem 90 and combine this with the bound from Theorem 52:

$$\begin{aligned}
\mathrm{ac}_k(B_p^m, B_q^m) &\overset{\text{Thm 90}}{\leq} \sum_{\ell \geq 1} \frac{1}{k2^\ell} \cdot \mathrm{vb}_{k2^\ell}(B_p^m, B_q^m) \\
&\overset{\text{Thm 52}}{\lesssim} \frac{1}{\frac{1}{2} - \frac{1}{p} + \frac{1}{q}} \sum_{\ell \geq 1} \frac{\sqrt{\min\{p, \log(\frac{2m}{k2^\ell})\}}}{(k2^\ell)^{1/2+1/p-1/q}} \\
&\lesssim \frac{1}{\frac{1}{2} - \frac{1}{p} + \frac{1}{q}} \cdot \frac{\sqrt{\min\{p, \log(\frac{2m}{k})\}}}{k^{1/2+1/p-1/q}}.
\end{aligned}$$

Note that the exponent $\alpha := 1/2 + 1/p - 1/q$ satisfies $\alpha \geq 1/2$ and so the sum is already dominated by the very first term. The running time to find the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ is dominated by $O(\log m)$ calls to find the coloring behind $\mathsf{vb}_t(B_p^m, B_q^m)$ where $t \leq m+1$ which results in a total running time of $O(\log^2 m) \cdot T(m, m+1) \leq O(m^{1+\omega} \log^2(m))$ as $\omega \geq 2$. $\square$

Remark 5. *In the case where $p = 2$ and $q = \infty$, we can apply Theorem 91 with $q' := \log_2 m$ to obtain $\mathrm{ac}_k(B_2^m, B_\infty^m) \lesssim \frac{\log m}{k}$ by noting that Lemma 94 implies $\|\boldsymbol{x}\|_\infty \leq \|\boldsymbol{x}\|_{\log_2 m} \leq 2\|\boldsymbol{x}\|_\infty$ for any $\boldsymbol{x} \in \mathbb{R}^m$. But for $P = B_2^m$ and $Q = B_\infty^m$ one has $\mathsf{vb}_k(B_2^m, B_\infty^m) \lesssim \sqrt{\log m}$ by the result of Banaszczyk [13], so that Theorem 90 yields the improved bound*

$$\mathrm{ac}_k(B_2^m, B_\infty^m) \lesssim \frac{\sqrt{\log m}}{k}.$$

It remains an interesting open question whether this may be improved to $O(\frac{1}{k})$.

Chapter 5

HEREDITARY NOTIONS OF DISCREPANCY

This chapter is based on a joint paper with Haotian Jiang [79].

5.1 Introduction

Given a matrix $A \in \mathbb{R}^{m\times n}$, the *discrepancy* of A is $\mathrm{disc}(A) := \min_{\boldsymbol{x}\in\{-1,+1\}^n} \|A\boldsymbol{x}\|_\infty$. The *hereditary discrepancy* of A is defined as $\mathsf{herdisc}(A) := \max_{S\subseteq[n]} \mathrm{disc}(A_S)$, where A_S denotes the restriction of the matrix A to columns in S. For a set system $\mathcal{F}$, $\mathrm{disc}(\mathcal{F})$ and $\mathsf{herdisc}(\mathcal{F})$ are defined to be $\mathrm{disc}(A_\mathcal{F})$ and $\mathsf{herdisc}(A_\mathcal{F})$, where $A_\mathcal{F}$ is the incidence matrix of $\mathcal{F}$.

In seminal work, Lovász, Spencer, and Vesztergombi [102] introduced a powerful tool, known as the *determinant lower bound*, for bounding hereditary discrepancy:

$$\mathsf{detLB}(A) := \max_{k\in\mathbb{N}} \max_{\substack{(S,T)\subseteq[m]\times[n]\\ |S|=|T|=k}} |\det(A_{S,T})|^{1/k},$$

where $A_{S,T}$ denotes the restriction of A to rows in S and columns in T. In particular, they showed that $\mathsf{herdisc}(A) \geq \frac{1}{2}\mathsf{detLB}(A)$ for any matrix A. A reverse relation was established by Matoušek [108], who showed that $\mathsf{herdisc}(A) \leq O(\log(mn)\sqrt{\log(n)}) \cdot \mathsf{detLB}(A)$. However, Matoušek's bound does not match the largest known gap of $\Theta(\log(n))$ between $\mathsf{herdisc}(A)$ and $\mathsf{detLB}(A)$, given by a construction of Pálvölgyi [128] or the counter-example to Beck's three permutation conjecture [123].

Our main result is the following improvement over Matoušek's bound in [108].

Theorem 98. *Given a matrix $A \in \mathbb{R}^{m\times n}$, one can efficiently find $\boldsymbol{x} \in \{+1,-1\}^n$ such that $\|A\boldsymbol{x}\|_\infty \leq O(\sqrt{\log(m)\cdot\log(n)} \cdot \mathsf{detLB}(A))$.*

Restricting to an arbitrary subset of the columns of A, one immediately obtains the following:

Corollary 99. *For any matrix $A \in \mathbb{R}^{m\times n}$, $\mathsf{herdisc}(A) \leq O(\sqrt{\log(m)\cdot\log(n)}\cdot \mathsf{detLB}(A))$.*

In light of the examples in [128, 123] where $\mathsf{herdisc}(A) \geq \Omega(\log n)\cdot\mathsf{detLB}(A)$, Theorem 98 is tight up to constants whenever $m = \mathsf{poly}(n)$. For the case where $m \gg \mathsf{poly}(n)$, one cannot hope to improve the $\sqrt{\log(m)}$ dependence on m in Theorem 98. In particular, the set system $\mathcal{F} = 2^{[n]}$ has $\mathsf{herdisc}(\mathcal{F}) = n$, $\mathsf{detLB}(\mathcal{F}) = \sqrt{n}$ and therefore $\mathsf{herdisc}(\mathcal{F}) \geq \sqrt{\log(m)}\cdot \mathsf{detLB}(\mathcal{F})$. Very recently, Li and Nikolov [101] showed that the upper bound is tight for almost all parameters m, n.

Hereditary discrepancy of union of set systems. A question of V. Sós (see [102]) asks whether $\mathsf{herdisc}(\mathcal{F}_1 \cup \mathcal{F}_2)$ can be estimated in terms of $\mathsf{herdisc}(\mathcal{F}_1)$ and $\mathsf{herdisc}(\mathcal{F}_2)$, for any set systems $\mathcal{F}_1$ and $\mathcal{F}_2$ over $[n]$. This is, however, not possible without any dependence on $m = |\mathcal{F}_1 \cup \mathcal{F}_2|$ or n, as first shown by an example of Hoffman (Proposition 4.11 in [111]). This can also be seen from the examples in [128, 123]. In [87], it was shown that $\mathsf{herdisc}(\mathcal{F}_1 \cup \mathcal{F}_2) \leq O(\log(n))\cdot \mathsf{herdisc}(\mathcal{F}_1)$ when $\mathcal{F}_2$ contains a single set. For more general set systems, Matoušek [108] proved that $\mathsf{herdisc}(\mathcal{F}) \leq O(\sqrt{t}\log(mn)\sqrt{\log(n)})\cdot \max_{i\in[t]}(\mathsf{herdisc}(\mathcal{F}_i))$, where $\mathcal{F} = \mathcal{F}_1 \cup \cdots \cup \mathcal{F}_t$ and $m = |\mathcal{F}|$.

Theorem 98 together with Lemma 4 in [108] immediately imply the following improvement of this result, whose proof is the same as in [108]. For $t = 2$ and $m = \mathsf{poly}(n)$, this bound is tight up to constants.

Theorem 100. *Let $\mathcal{F}$ be a system of m sets on $[n]$ such that $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2 \cup \cdots \cup \mathcal{F}_t$. Then,*

$$\mathsf{herdisc}(\mathcal{F}) \leq O\left(\sqrt{t\log(m)\log(n)}\right)\cdot \max_{i\in[t]}(\mathsf{herdisc}(\mathcal{F}_i)).$$

Approximating hereditary discrepancy. It was shown in [44] that $\mathrm{disc}(A)$ cannot be approximated in polynomial time for an arbitrary matrix $A \in \{0,1\}^{m\times n}$. The more robust notion of hereditary discrepancy, however, can be approximated within a polylog factor. The best-known result in this direction is a $O(\log(\min(m,n))\cdot\sqrt{\log(m)})$-approximation to hereditary discrepancy via the γ_2-norm [109]. When $m = \mathsf{poly}(n)$, this approximation factor is $O(\log^{3/2}(n))$.

Our result in Theorem 98 suggests a potential approach of approximating hereditary discrepancy by approximating the determinant lower bound. There has been a recent line of work in approximating the maximum $k \times k$ subdeterminant for a given matrix A. For $k = \min(m, n)$, Nikolov [124] gave a $2^{O(k)}$-approximation; for general values of k, Anari and Vuong [6] showed a $k^{O(k)}$-approximation algorithm. If these results can be strengthened to a $2^{O(k)}$-approximation algorithm for general values of k, then together with Theorem 98, one would obtain the first $O(\log(n))$-approximation algorithm for hereditary discrepancy when $m = \mathsf{poly}(n)$.

Overview of proof of Theorem 98. We follow the approaches in [16] and [108]. The key notion to prove Theorem 98 is that of *hereditary partial vector discrepancy*, which is defined as follows. Given a matrix $A \in \mathbb{R}^{m \times n}$ with entries a_{ij} for $i \in [m]$ and $j \in [n]$, we consider the following SDP for a subset $S \subseteq [n]$ and a parameter $\lambda \geq 0$:

$$\begin{aligned} \Big\| \sum_{j \in S} a_{ij} \boldsymbol{v}_j \Big\|_2^2 &\leq \lambda^2 \quad \forall i \in [m], \\ \sum_{j=1}^{n} \|\boldsymbol{v}_j\|_2^2 &\geq |S|/2, \\ \|\boldsymbol{v}_j\|_2^2 &\leq 1 \quad \forall j \in S, \\ \|\boldsymbol{v}_j\|_2^2 &= 0 \quad \forall j \in [n] \setminus S. \end{aligned}$$

Define the *partial vector discrepancy* of A, denoted as $\mathsf{pvdisc}(A)$, to be the smallest value of λ such that the above SDP is feasible, and *hereditary partial vector discrepancy* $\mathsf{herpvdisc}(A)$ to be the smallest λ such that $\mathsf{SDP}(A, S, \lambda)$ is feasible for any subset $S \subseteq [n]$.

Using the above definition, we show in Lemma 102 of Section 5.2.1 that the above SDP can be rounded efficiently to obtain a coloring with discrepancy at most $O(\sqrt{\log(m)\log(n)} \cdot \mathsf{herpvdisc}(A))$. We then prove in Lemma 104 of Section 5.2.2 that $\mathsf{herpvdisc}(A) \leq O(\mathsf{detLB}(A))$, from which Theorem 98 immediately follows. We conjecture that $\mathsf{herpvdisc}(A)$ is the same as $\mathsf{detLB}(A)$ up to constants (Conjecture 13).

Notations and preliminaries. Given a matrix $A \in \mathbb{R}^{m \times n}$, its rows will be denoted by

$\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m \in \mathbb{R}^n$. Define $A_{S,T}$ to be the matrix with rows restricted to some subset $S \subseteq [m]$ and columns restricted to some $T \subseteq [n]$, and $A_S := A_{[m],S}$.

Theorem 101 (Freedman's Inequality, Theorem 1.6 in [66])**.** *Consider a real-valued martingale sequence $\{X_t\}_{t\geq 0}$ such that $X_0 = 0$, and $\mathbb{E}[X_{t+1}|\mathcal{F}_t] = 0$ for all t, where $\{\mathcal{F}_t\}_{t\geq 0}$ is the filtration defined by the martingale. Assume that the sequence is uniformly bounded, i.e., $|X_t| \leq M$ almost surely for all t. Now define the predictable quadratic variation process of the martingale to be $W_t = \sum_{j=1}^t \mathbb{E}[X_j^2|\mathcal{F}_{j-1}]$ for all $t \geq 1$. Then for all $\ell \geq 0$ and $\sigma^2 > 0$ and any stopping time τ, we have*

$$\Pr\Big[\Big|\sum_{j=0}^{\tau} X_j\Big| \geq \ell \wedge W_\tau \leq \sigma^2 \textit{for some stopping time } \tau\Big] \leq 2\exp\Big(-\frac{\ell^2/2}{\sigma^2 + M\ell/3}\Big).$$

5.2 *Proof of Theorem 98*

5.2.1 The Algorithm

The main result of this subsection is the following lemma.

Lemma 102. *Given a matrix $A \in \mathbb{R}^{m\times n}$, there exists a randomized algorithm that w.h.p. constructs a coloring $\boldsymbol{x} \in \{+1,-1\}^n$ such that $\|A\boldsymbol{x}\|_\infty \leq O(\sqrt{\log(m)\log(n)} \cdot \mathsf{herpvdisc}(A))$. This implies that $\mathsf{herdisc}(A) \leq O(\sqrt{\log(m)\log(n)} \cdot \mathsf{herpvdisc}(A))$.*

The algorithm in Lemma 102 is given in Algorithm 1. This algorithm is a variant of the random walk in [16], using the SDP for hereditary partial vector discrepancy.

Since Lemma 102 is invariant under rescaling of the matrix A, we may assume without loss of generality that that $\max_{i,j} |a_{i,j}| = 1$. Given a coloring $\boldsymbol{x} \in [-1,1]^n$, we say an element $i \in [n]$ is alive if $|x(i)| < 1 - 1/n$. The following lemma from [16] states that the number of alive elements halves after $O(1/s^2)$ steps.

Lemma 103 (Lemma 4.1 of [16])**.** *Let $\boldsymbol{y} \in [-1,+1]^n$ be an arbitrary fractional coloring with at most k alive variables. Let $\boldsymbol{z}$ be the fractional coloring obtained by running algorithm 1 with $\boldsymbol{x}_0' = \boldsymbol{y}$ for $T' = 16/s^2$ steps. Then the probability that $\boldsymbol{z}$ has at least $k/2$ alive variables is at most* $1/4$.

Algorithm 1 $\textsc{HerpvdiscRounding}(A)$

$\lambda \leftarrow \mathsf{herpvdisc}(A)$ ▷ The value of λ can be approximated with a binary search
$\boldsymbol{x}_0 \leftarrow \mathbf{0} \in \mathbb{R}^n$, $S_0 \leftarrow [n]$, $s \leftarrow 1/m^2n^2$, $T \leftarrow 200\log(n)/s^2$
for $t = 1, 2, \cdots, T$ **do**
  $\boldsymbol{v}_1, \cdots, \boldsymbol{v}_n \leftarrow \mathsf{SDP}(A, S_{t-1}, \lambda)$
  Sample $\boldsymbol{r} \in \{-1, +1\}^n$ uniformly at random
  for $i \in [n]$ **do** $x_t(i) \leftarrow x_{t-1}(i) + s \cdot \langle \boldsymbol{r}, \boldsymbol{v}_i \rangle$
  end for
  $S_t \leftarrow S_{t-1}$
  for $i \in [n] \setminus S_{t-1}$ **do**
    if $|x_t(i)| \geq 1 - 1/n$ **then**
      $S_t \leftarrow S_t \setminus \{i\}$
    end if
  end for
end for
Round $\boldsymbol{x}_T$ to a vector $\boldsymbol{x} \in \{-1, +1\}^n$
Return $\boldsymbol{x}$

Proof of Lemma 102. We first argue that after $T = 400\log(n)/s^2$ steps, no element is alive with high probability. Divide the time horizon into epochs of size $16/s^2$. For each epoch, Lemma 103 states that regardless of the past, the number of alive elements decreases by at least half with probability at least $3/4$. It follows that no element is alive with high probability after $25\log(n)$ epochs. Note that when no element is alive for the coloring $\boldsymbol{x}_T$, one can round it to a full coloring without changing the discrepancy of each set by more than 1.

Next we prove that with high probability, the discrepancy of each row of A is at most $O(\sqrt{\log(m)\log(n)}) \cdot \lambda$. We consider any $j \in [m]$, and denote $\mathrm{disc}_t(j) = \langle \boldsymbol{a}_j, \boldsymbol{x}_t \rangle$ the discrepancy of row j at the end of time step $t \in [T]$. Note that $\mathbb{E}[\mathrm{disc}_t(j) - \mathrm{disc}_{t-1}(j) | \mathrm{disc}_{t-1}(j)] = 0$ and $\mathbb{E}[(\mathrm{disc}_t(j) - \mathrm{disc}_{t-1}(j))^2 | \mathrm{disc}_{t-1}(j)] \leq \lambda^2 s^2$. It follows from Freedman's inequality (Theorem 101) that

$$\Pr\left[|\mathrm{disc}_T(j)| \geq 10\sqrt{\log(m)\log(n)} \cdot \lambda \right] \leq 1/m^2.$$

So by the union bound, the discrepancy of the obtained coloring is at most $O(\sqrt{\log(m)\log(n)} \cdot \mathsf{herpvdisc}(A))$ with high probability. This completes the proof of Lemma 102. $\square$

5.2.2 *Bounding Partial Vector Discrepancy*

In this subsection, we prove the following lemma which upper bounds partial vector discrepancy in terms of the determinant lower bound. The proof can be seen as a simplification of Lemma 8 in [108], which gives a corresponding upper bound for *vector discrepancy* that is weaker by a factor of $\sqrt{\log n}$ due to a bucketing argument that is not needed here.

Lemma 104. *For any $A \in \mathbb{R}^{m \times n}$, we have* $\mathsf{herpvdisc}(A) \leq O(\mathsf{detLB}(A))$.

Proof. Recall that $\mathsf{pvdisc}(A)^2$ is the optimal value of the SDP given by

$$\begin{aligned} \min\ & t \\ \Big\| \sum_{j=1}^{n} a_{ij} \boldsymbol{v}_j \Big\|_2^2 \ &\leq\ t \quad \forall i \in [m] \\ \sum_{j=1}^{n} \|\boldsymbol{v}_j\|_2^2 \ &\geq\ n/2 \\ \|\boldsymbol{v}_j\|_2^2 \ &\leq\ 1 \quad \forall j \in [n]. \end{aligned}$$

By denoting $X_{ij} := \langle \boldsymbol{v}_i, \boldsymbol{v}_j \rangle$, we may rewrite this SDP as follows:

$$\begin{aligned} \min\ & t \\ \langle \boldsymbol{a}_i \boldsymbol{a}_i^\top, X \rangle \ &\leq\ t \quad \forall i \in [m] \\ \langle I_n, X \rangle \ &\geq\ n/2 \\ \langle \boldsymbol{e}_j \boldsymbol{e}_j^\top, X \rangle \ &\leq\ 1 \quad \forall j \in [n] \\ X \ &\succeq\ 0, \end{aligned}$$

where $\boldsymbol{e}_j$ denotes the vector with 1 on the j-th coordinate and 0 elsewhere. The dual formulation of the above SDP is given by the following:

$$\begin{aligned} \max \quad & n\gamma - \sum_{j=1}^{n} z_j \\ \sum_{i=1}^{m} w_i \boldsymbol{a}_i \boldsymbol{a}_i^\top + \sum_{j=1}^{n} z_j \boldsymbol{e}_j \boldsymbol{e}_j^\top \ &\succeq\ 2\gamma \cdot I_n \\ \sum_{i=1}^{m} w_i \ &=\ 1 \\ \boldsymbol{w}, \boldsymbol{z} \ &\geq\ 0. \end{aligned}$$

Denote $\lambda := \mathsf{pvdisc}(A)$. By Slater's condition, there exists a feasible dual solution $(\boldsymbol{w}, \boldsymbol{z}, \gamma)$ such that $\boldsymbol{w}, \boldsymbol{z} \geq 0$ and $n\gamma - \sum_{j=1}^{n} z_j = \lambda^2$. Indeed, the dual has a feasible interior point (for example, $w_i = 1/m, z_j = 1$ and $\gamma = 0$) and is bounded, since we may rewrite the first constraint as

$$\sum_{i=1}^{m} w_i \boldsymbol{a}_i \boldsymbol{a}_i^\top \succeq \sum_{j=1}^{n} (2\gamma - z_j) \cdot \boldsymbol{e}_j \boldsymbol{e}_j^\top, \tag{5.1}$$

which implies

$$n\gamma - \sum_{j=1}^{n} z_j \leq n\gamma - \frac{1}{2}\sum_{j=1}^{n} z_j \leq \frac{1}{2}\mathrm{Tr}\Big[\sum_{i=1}^{m} \boldsymbol{a}_i\boldsymbol{a}_i^\top\Big].$$

Let $\tilde{A}$ be the matrix obtained from A by multiplying the i-th row by $\sqrt{w_i}$ and $J \subseteq [n]$ be the set of columns for which $z_j < \frac{3}{2}\gamma$. Note that $|J| \geq \frac{1}{3}n$, for otherwise $\sum_{j=1}^{n} z_j > \frac{2}{3}n \cdot \frac{3}{2}\gamma = n\gamma$. Since for each $j \in J$ we have $2\gamma - z_j \geq \frac{1}{2}\gamma$, for any vector $\boldsymbol{x} \in \mathbb{R}^J$ it follows by (5.1):

$$\boldsymbol{x}^\top \widetilde{A}_J^\top \widetilde{A}_J \boldsymbol{x} \geq \frac{1}{2}\gamma \cdot \|\boldsymbol{x}\|_2^2 \geq \frac{\lambda^2}{2n} \cdot \|\boldsymbol{x}\|_2^2.$$

This implies that all eigenvalues of $\widetilde{A}_J^\top \widetilde{A}_J$ are at least $\lambda^2/2n$, so that $\det(\widetilde{A}_J^\top \widetilde{A}_J) \geq (\lambda^2/2n)^{|J|}$. In the other direction, the Cauchy-Binet formula also gives

$$\begin{aligned}
\det(\widetilde{A}_J^\top \widetilde{A}_J) &= \sum_{\substack{I \subseteq [m] \\ |I|=|J|}} \det(\widetilde{A}_{I,J})^2 \\
&= \sum_{\substack{I \subseteq [m] \\ |I|=|J|}} \det(A_{I,J})^2 \prod_{i \in I} w_i \\
&\leq \mathsf{detLB}(A)^{2|J|} \cdot \sum_{\substack{I \subseteq [m] \\ |I|=|J|}} \prod_{i \in I} w_i \\
&\leq \mathsf{detLB}(A)^{2|J|} \cdot \frac{1}{|J|!}\Big(\sum_{i=1}^{m} w_i\Big)^{|J|},
\end{aligned}$$

where the last inequality follows as each term $\prod_{i \in I} w_i$ appears $|J|!$ times in $\Big(\sum_{i=1}^{m} w_i\Big)^{|J|}$. Since $\sum_{i=1}^{m} w_i = 1$, we conclude

$$\mathsf{detLB}(A)^{2|J|} \cdot \frac{1}{|J|!} \geq \det(\widetilde{A}_J^\top \widetilde{A}_J) \geq (\lambda^2/2n)^{|J|},$$

from which $\mathsf{detLB}(A) \geq \Omega(\lambda \cdot \sqrt{|J|/n}) = \Omega(\lambda) = \Omega(\mathsf{pvdisc}(A))$. Applying this result to all subsets $S \subseteq [n]$ of the columns of A proves the lemma. $\square$

We conjecture that the above Lemma 104 is tight up to constants.

Conjecture 13. *For any matrix $A \in \mathbb{R}^{m \times n}$, we have* $\mathsf{detLB}(A) = \Theta(\mathsf{herpvdisc}(A))$.

Chapter 6

MATRIX BALANCING I: GRAPH SPARSIFICATION

This chapter is based on joint work with Thomas Rothvoss [138] and with Arun Jambulapati and Kevin Tian [77].

6.1 Introduction

Discrepancy theory is a subfield of combinatorics with several applications to theoretical computer science, see for example the books [111, 45]. In the classical setting one is given a family of sets $\mathcal{S} = \{S_1, \dots, S_m\}$ with $S_i \subseteq \{1, \dots, n\}$ and the goal is to find a *coloring* $\chi : [n] \to \{-1, +1\}$ so that the maximum imbalance $\max_{S \in \mathcal{S}} |\sum_{j \in S} \chi(j)|$ is minimized. This minimum value is called the *discrepancy* of the family, denoted by disc($\mathcal{S}$). A seminal result of Spencer [157] says that for any set family one has disc($\mathcal{S}$) $\leq O(\sqrt{n \log(2m/n)})$, assuming that $m \geq n$. It is instructive to observe that for $m = n$, Spencer's result gives the bound of $O(\sqrt{n})$, while a uniform random coloring will have a discrepancy of $O(\sqrt{n \log(n)})$. Moreover, one can show that for some set systems, only an exponentially small fraction of all colorings will indeed have a discrepancy of $O(\sqrt{n})$. This demonstrates that in fact, Spencer's result provides the existence of a rather rare object.

The cleanest approach to prove Spencer's result is due to Giannopoulos [67], which we sketch for $m = n$: Consider the set $K = \{\boldsymbol{x} \in \mathbb{R}^n : |\sum_{j \in S_i} x_j| \leq \sqrt{n} \; \forall i \in [n]\} = \bigcap_{i \in [n]} Q_i$, a symmetric convex body which denotes the set of good-enough fractional colorings. Here Q_i is the strip of colorings that are good for set S_i. The Lemma of Sidak-Khatri [86, 155] allows us to lower bound the *Gaussian measure* of K as $\gamma_n(K) \geq \prod_{i=1}^n \gamma_n(Q_i) \geq e^{-cn}$ for some constant $c > 0$ using that each strip Q_i has a constant width. This rather weak bound on the measure is sufficient to use a *pigeonhole principle* argument and conclude that $c'K$ must contain a partial coloring $\boldsymbol{x} \in \{-1,0,1\}^n$ with $|\text{supp}(\boldsymbol{x})| \geq \frac{n}{2}$. Then one can color the elements in supp($\boldsymbol{x}$) accordingly and repeat the argument for the remaining uncolored elements. The overall $O(\sqrt{n})$ bound follows from the fact that the discrepancy of the partial colorings decreases geometrically as the number of elements in the set system decreases.

While the pigeonhole principle based argument above is non-constructive in nature, Bansal [16] designed a polynomial time algorithm for finding the coloring guaranteed by Spencer's Theorem. Here, [16] exploits that it suffices to obtain a good enough *fractional* partial coloring $\boldsymbol{x} \in [-1,1]^n$ with a constant fraction of entries in $\{-1,1\}$ to make the argument work. Later, Lovett and Meka [103] found a Brownian motion-type algorithm that — despite being a lot simpler — works for more general polyhedral settings. Finally, the random projection algorithm of Rothvoss [146] works for arbitrary symmetric convex bodies that satisfy the measure lower bound. Another remarkable result is due to Bansal, Dadush, Garg and Lovett [19]: for any symmetric body K with $\gamma_n(K) \geq \frac{1}{2}$ and any vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m \in \mathbb{R}^n$ of length $\|\boldsymbol{v}_i\|_2 \leq 1$, one can find signs $\boldsymbol{x} \in \{-1,1\}^m$ in randomized polynomial time so that $\sum_{i=1}^m x_i \boldsymbol{v}_i \in O(1) \cdot K$. This was known before by a non-constructive convex geometric argument due to Banaszczyk [13].

There are two possible strengthenings of Spencer's Theorem that are both open at the time of this writing: suppose that the set system is *sparse* in the sense that every element is in at most t sets. It is known that disc($\mathcal{S}$) $\leq 2t$ [33] as well as disc($\mathcal{S}$) $\leq O(\sqrt{t\log(n)})$ [13, 19], while the Beck-Fiala Conjecture suggests that disc($\mathcal{S}$) $\leq O(\sqrt{t})$ is the right bound. For the second generalization — the one that we are following in this paper — it is helpful to define $\boldsymbol{A}_i$ as the $m \times m$ diagonal matrix with (j,j) entry 1 if $i \in S_j$ and 0 otherwise. If $\|\cdot\|_{\text{op}}$ denotes the maximum singular value of a matrix, then Spencer's result can be interpreted as the existence of a coloring $\boldsymbol{x} \in \{-1,1\}^n$ so that $\|\sum_{i=1}^n x_i \boldsymbol{A}_i\|_{\text{op}} \leq O(\sqrt{n\log(2m/n)})$. A conjecture raised by Meka[1] is whether for $m = n$,

[1]See the blog post `https://windowsontheory.org/2014/02/07/discrepancy-and-beating-the-union-bound/`

this bound is also possible for arbitrary symmetric matrices $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_n \in \mathbb{R}^{n \times n}$ that satisfy $\|\boldsymbol{A}_i\|_{\text{op}} \leq 1$. One can prove using *matrix concentration inequalities* that a random coloring $\boldsymbol{x}$ will lead to $\| \sum_{i=1}^n x_i \boldsymbol{A}_i \|_{\text{op}} \leq O(\sqrt{n \log(n)})$, and the same bound can also be achieved deterministically using a matrix multiplicative weight update argument [171]. An excellent overview of matrix concentration can be found in the monograph of Tropp [164].

To understand the difficulty of proving Meka's conjecture, assume $m = n$ and revisit the approach of Giannopoulos for Spencer's Theorem. We can again define a set

$$K := \Big\{ \boldsymbol{x} \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i \boldsymbol{A}_i \Big\|_{\text{op}} \leq \sqrt{n} \Big\} = \Big\{ \boldsymbol{x} \in \mathbb{R}^n : \sum_{i=1}^n x_i \left\langle \boldsymbol{A}_i, \boldsymbol{y}\boldsymbol{y}^\top \right\rangle \leq \sqrt{n} \;\; \forall \boldsymbol{y} \in \mathbb{R}^m : \|\boldsymbol{y}\|_2 = 1 \Big\}$$

of good enough fractional colorings. Since $\| \cdot \|_{\text{op}}$ is a norm, K will indeed be symmetric and convex. It would hence suffice to prove that $\gamma_n(K) \geq 2^{-cn}$ for some constant $c > 0$. However, it is open whether this inequality holds. The issue is that K is non-polyhedral and applying Sidak-Khatri's bound over infinitely[2] many vectors $\boldsymbol{y}$ is way too inefficient. While matrix concentration inequalities are fantastic at proving that likely events are indeed likely, they seem to be unable to prove that unlikely events are not too unlikely. With a scaling argument, they can still be used to prove that $\gamma_n(K) \geq (\log(n))^{-cn}$ for some constant $c > 0$, assuming $m = n$, though better bounds seem out of reach.

In terms of discrepancy in spectral settings, a different line of techniques has been arguably more successful. A beautiful and influential paper by Batson, Spielman and Srivastava [31] proves that for any undirected graph on n nodes one can take a weighted subgraph with just a *linear* number of edges that approximates every cut within a constant factor. Translated into linear algebra terms, [31] show that given any vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m \in \mathbb{R}^n$ that are in *isotropic position*, i.e. $\sum_{i=1}^m \boldsymbol{v}_i \boldsymbol{v}_i^\top = \boldsymbol{I}_n$, one can find weights $\boldsymbol{s} \in \mathbb{R}^m_{\geq 0}$ with $|\text{supp}(\boldsymbol{s})| \leq O(n/\varepsilon^2)$ so that $(1-\varepsilon) \cdot \boldsymbol{I}_n \preceq \sum_{i=1}^m s_i \boldsymbol{v}_i \boldsymbol{v}_i^\top \preceq (1+\varepsilon) \cdot \boldsymbol{I}_n$, and indeed Lee and Sun showed this can be done in nearly linear time [96]. In a more recent celebrated paper, Marcus, Spielman and Srivastava [106] resolved the *Kadison-Singer Conjecture*, a problem that

[2] One can use an ε-net of $2^{\Theta(m)}$ many vectors $\boldsymbol{y}$ but the bound is still too weak.

has appeared independently in different forms in many areas of mathematics. In a simple-to-state version, their result says that for any vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_m \in \mathbb{R}^n$ with $\sum_{i=1}^m \boldsymbol{v}_i \boldsymbol{v}_i^\top = \boldsymbol{I}_n$ and $\|\boldsymbol{v}_i\|_2 \leq \varepsilon$ for all $i \in [m]$, there are signs $\boldsymbol{x} \in \{-1, 1\}^m$ so that $\| \sum_{i=1}^m x_i \boldsymbol{v}_i \boldsymbol{v}_i^\top \|_{\text{op}} \leq O(\varepsilon)$. On a very high level view, both methods of [31] and [106] control a carefully chosen potential function, though we note there is still no known polynomial time algorithm for the latter.

The goal of this paper will be to connect the classical discrepancy theory and the spectral discrepancy theory of [31, 106] and develop arguments that prove largeness of non-polyhedral bodies. We remark that we made no attempt at optimizing constants but rather prefer to keep the exposition simple.

Notation. For a (not necessarily symmetric) matrix $\boldsymbol{M} \in \mathbb{R}^{n \times n}$ the operator norm can be formally defined as $\|\boldsymbol{M}\|_{\text{op}} := \max\{\|\boldsymbol{M}\boldsymbol{x}\|_2 : \boldsymbol{x} \in \mathbb{R}^n \text{ with } \|\boldsymbol{x}\|_2 = 1\}$. For a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ with eigendecomposition $\boldsymbol{A} = \sum_{i=1}^n \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$, we write $|\boldsymbol{A}| := \sum_{i=1}^n |\lambda_i| \boldsymbol{v}_i \boldsymbol{v}_i^\top$ as the matrix where all eigenvalues have been replaced by their absolute values. In this notation, $\|\boldsymbol{A}\|_{\text{op}} := \max\{|\lambda_i| : i \in [n]\}$ is the maximum singular value. We abbreviate $B_2^n := \{\boldsymbol{x} \in \mathbb{R}^n \mid \|\boldsymbol{x}\|_2 \leq 1\}$ and $S^{n-1} := \{\boldsymbol{x} \in \mathbb{R}^n \mid \|\boldsymbol{x}\|_2 = 1\}$. Given symmetric matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n \times n}$, we write $\boldsymbol{A} \preceq \boldsymbol{B}$ if $\boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x} \leq \boldsymbol{x}^\top \boldsymbol{B} \boldsymbol{x}$ for all $\boldsymbol{x} \in \mathbb{R}^n$.

A *convex body* is a closed convex set $K \subset \mathbb{R}^n$ with nonempty interior. We denote $d(\boldsymbol{y}, K) := \min_{\boldsymbol{x} \in K} \|\boldsymbol{x} - \boldsymbol{y}\|_2$ as the *distance* from $\boldsymbol{y}$ to K. Let $K_\delta = \{\boldsymbol{x} \in \mathbb{R}^n \mid d(\boldsymbol{x}, K) \leq \delta\}$ be the set of points that have distance at most δ to K (in particular, $K \subseteq K_\delta$). The *Minkowski sum* of sets A and B is defined as $A + B := \{\boldsymbol{a} + \boldsymbol{b} \mid \boldsymbol{a} \in A, \boldsymbol{b} \in B\}$. A *halfspace* is a set of the form $H := \{\boldsymbol{x} \in \mathbb{R}^n | \langle \boldsymbol{v}, \boldsymbol{x} \rangle \leq \lambda\}$ for some $\boldsymbol{v} \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$. The *Gaussian measure* of K is defined as $\gamma_n(K) := \Pr_{\boldsymbol{y} \sim N(\boldsymbol{0}, \boldsymbol{I}_n)}[\boldsymbol{y} \in K]$. Here $N(\boldsymbol{0}, \boldsymbol{I}_n)$ is the distribution of a standard Gaussian in $\mathbb{R}^n$.

6.1.1 *Our contribution*

A possible way to approach the setting of Batson, Spielman, Srivastava [31] from a classical discrepancy perspective is to take vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m$ in isotropic position and consider the body $K = \{\boldsymbol{x} \in \mathbb{R}^m \mid \| \sum_{i=1}^m x_i \boldsymbol{v}_i \boldsymbol{v}_i^\top \|_{\text{op}} \leq \sqrt{n/m}\}$. If we could prove that $\gamma_m(K) \geq 2^{-cm}$, then the algorithm of [146] would be able to find a partial coloring. We first prove an expansion measure lower bound in a slightly more general setting:

Theorem 105. *Let $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m \in \mathbb{R}^{n \times n}$ be $m \geq \max\{200, 6\log n\}$ symmetric matrices and let $\boldsymbol{S} := \sum_{i=1}^m |\boldsymbol{A}_i|$ satisfy $\boldsymbol{S} \preceq \boldsymbol{I}_n$. Denote $\tau := \mathrm{Tr}[\boldsymbol{S}]$ and $\varepsilon := \sqrt{\tau/m}$. Suppose that $\tau \geq 1$. Then for any $0 < \alpha < 1$, the set*

$$K := \Big\{ \boldsymbol{x} \in \mathbb{R}^m \mid \Big\| \sum_{i=1}^m x_i \boldsymbol{A}_i \Big\|_{op} \leq \varepsilon \Big\}$$

satisfies $\gamma_m\Big(\frac{50}{\alpha} K + \alpha\sqrt{m} B_2^m\Big) \geq \frac{1}{2}$. That is, $\Pr_{\boldsymbol{y} \sim N(\boldsymbol{0}, \boldsymbol{I}_m)} \Big[d\Big(\boldsymbol{y}, \frac{50}{\alpha} K\Big) \leq \alpha\sqrt{m} \Big] \geq \frac{1}{2}$.

Note that in particular the rank-1 matrices $\boldsymbol{A}_i = \boldsymbol{v}_i \boldsymbol{v}_i^T$ with $\sum_{i=1}^m \boldsymbol{v}_i \boldsymbol{v}_i^T \preceq \boldsymbol{I}_n$ satisfy the premise of Theorem 105. By Theorem 10 we get the desired measure lower bound on K:

Theorem 106. *Let $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m \in \mathbb{R}^{n \times n}$ be $m \geq \max\{200, 6\log n\}$ symmetric matrices and let $\boldsymbol{S} := \sum_{i=1}^m |\boldsymbol{A}_i|$ satisfy $\boldsymbol{S} \preceq \boldsymbol{I}_n$. Denote $\tau := \mathrm{Tr}[\boldsymbol{S}]$ and $\varepsilon := \sqrt{\tau/m}$. Suppose that $\tau \geq 1$. Then for any $0 < \alpha < 1$, the set*

$$K := \Big\{ \boldsymbol{x} \in \mathbb{R}^m \mid \Big\| \sum_{i=1}^m x_i \boldsymbol{A}_i \Big\|_{op} \leq \varepsilon \Big\}$$

satisfies $\gamma_m(K) \geq 2^{-cm}$ for some universal constant $c > 0$.

A rather immediate consequence of this insight is that the following sampling algorithm will work with very high probability:

Spectral Sparsification Algorithm
• **Input:** PSD matrices $\boldsymbol{A}_1, \dots, \boldsymbol{A}_m \in \mathbb{R}^{n\times n}$ with $\sum_{i=1}^m \boldsymbol{A}_i = \boldsymbol{I}_n$ and $\varepsilon > 0$
• **Output:** $\boldsymbol{s} \in \mathbb{R}^m_{\geq 0} : \lvert\text{supp}(\boldsymbol{s})\rvert \leq \frac{n}{\varepsilon^2}$ and $(1-O(\varepsilon))\boldsymbol{I}_n \preceq \sum_{i=1}^m s_i \boldsymbol{A}_i \preceq (1+O(\varepsilon))\boldsymbol{I}_n$
(1) Set $s_i := 1$ for $i \in [m]$
(2) WHILE $\lvert\text{supp}(\boldsymbol{s})\rvert > \frac{n}{\varepsilon^2}$ DO
(3) Let $K := \{\boldsymbol{x} \in \mathbb{R}^{\text{supp}(\boldsymbol{s})} : \lVert \sum_{i\in\text{supp}(\boldsymbol{s})} x_i s_i \boldsymbol{A}_i \rVert_{\text{op}} \leq 1000\tilde{\varepsilon}\}$ with $\lvert\text{supp}(\boldsymbol{s})\rvert = \frac{n}{\tilde{\varepsilon}^2}$.
(4) Draw a Gaussian $\boldsymbol{y}^* \sim N(\boldsymbol{0}, \boldsymbol{I}_{\text{supp}(\boldsymbol{s})})$.
(5) Compute $\boldsymbol{x}^* := \text{argmin}\{\lVert \boldsymbol{x} - \boldsymbol{y}^*\rVert_2 : \boldsymbol{x} \in [-1,1]^{\text{supp}(s)} \cap K\}$.
(6) If $\#(i : x_i^* = -1) < \#(i : x_i = 1)$ then replace $\boldsymbol{x}^*$ by $-\boldsymbol{x}^*$.
(7) Update $s_i := s_i \cdot (1 + x_i^*)$.

In fact we will prove:

Theorem 107. *With probability at least* $1 - 2^{-\Omega(n)}$ *a run of the* Spectral Sparsification Algorithm *satisfies all of the following properties:* (a) *the algorithm runs in polynomial time;* (b) *the while loop is iterated at most* $O(\log m)$ *times;* (c) *at the end one has* $\lvert\text{supp}(\boldsymbol{s})\rvert \leq \frac{n}{\varepsilon^2}$ *and* $(1-O(\varepsilon))\boldsymbol{I}_n \preceq \sum_{i=1}^m s_i \boldsymbol{A}_i \preceq (1+O(\varepsilon))\boldsymbol{I}_n$.

Note that our algorithm produces sparse vector $\boldsymbol{s}$ by iteratively finding low discrepancy colorings. This technique has appeared before in the literature. For example for a set system with bounded *VC dimension*, one can prove the existence of small ε*-nets* in this manner. We refer to Chapter 4 of Chazelle's book [45] for details.

6.2 Preliminaries

In this section, we discuss several tools from probability and linear algebra that we will be using in the proofs.

Concentration. We need two concentration inequalities. For the first one, see [166].

Theorem 108. *If* $F : \mathbb{R}^m \to \mathbb{R}$ *is* 1*-Lipschitz, then for* $t \geq 0$ *one has*

$$\Pr_{\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{I}_m)}[F(\boldsymbol{y}) > \mathbb{E}[F(\boldsymbol{y})] + t] \leq e^{-t^2/2}.$$

For the proof of the following Corollary, see Section 6.5.

Corollary 109. *For $m \geq 7$ we have*

$$\Pr_{\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{I}_m)}\Big[\|\boldsymbol{y}\|_2 > m\Big] \leq 2^{-m} \quad \textit{and} \quad \mathop{\mathbb{E}}_{\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{I}_m)}\Big[\|\boldsymbol{y}\|_2^2 \mid \|\boldsymbol{y}\| \leq m\Big] \geq (1-2^{-m}) \cdot m.$$

We also need Azuma's inequality for Martingales with bounded increments, see [4].

Theorem 110 (**Azuma's Inequality**). *Let $0 = X_0, \dots, X_T$ be a Martingale with $|X_t - X_{t-1}| \leq a$ for all $t = 1, \dots, T$. Then for any $\lambda \geq 0$ we have*

$$\Pr[X_T > \lambda\sqrt{T}] \leq e^{-\lambda^2/2a^2}$$

Gaussians. In order to increase the measure from $\frac{1}{2}$ to $1 - 2^{-\Omega(m)}$ we use the following key theorem, see [95].

Theorem 111 (**Gaussian Isoperimetric Inequality**). *Let $K \subset \mathbb{R}^n$ be a measurable set and H be a halfspace such that $\gamma_n(K) = \gamma_n(H)$. Then $\gamma_n(K_\delta) \geq \gamma_n(H_\delta)$ for all $\delta > 0$.*

The following simple result is useful for dealing with dilations, see [163].

Theorem 112. *Let $K \subset \mathbb{R}^n$ be a measurable set and B be a closed Euclidean ball such that $\gamma_n(K) = \gamma_n(B)$. Then $\gamma_n(tK) \geq \gamma_n(tB)$ for all $t \in [0,1]$.*

For (not neccesarily symmetric) matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n\times n}$ we define the *Frobenius inner product* $\langle \boldsymbol{A}, \boldsymbol{B}\rangle_F := \mathrm{Tr}[\boldsymbol{A}^\top \boldsymbol{B}] = \sum_{i=1}^n \sum_{j=1}^n A_{ij}B_{ij}$ and the corresponding *Frobenius norm* $\|\boldsymbol{A}\|_F := \sqrt{\langle \boldsymbol{A}, \boldsymbol{A}\rangle_F} = (\sum_{i=1}^n \sum_{j=1}^n A_{ij}^2)^{1/2}$. Generalizing earlier notation, for a PSD matrix $\boldsymbol{X} \in \mathbb{R}^{m\times m}$, we define $N(\boldsymbol{0}, \boldsymbol{X})$ as the distribution of a centered Gaussian with covariance matrix $\boldsymbol{X}$. Note that there is a canonical way to generate such a distribution: let $X_{ij} = \langle \boldsymbol{v}_i, \boldsymbol{v}_j\rangle$ be the factorization of that matrix for some vectors $\boldsymbol{v}_i \in \mathbb{R}^r$. Then draw a standard Gaussian $\boldsymbol{y} \sim N(\boldsymbol{0}, \boldsymbol{I}_r)$, so that $(\langle \boldsymbol{g}, \boldsymbol{v}_1\rangle, \dots, \langle \boldsymbol{g}, \boldsymbol{v}_m\rangle) \sim N(\boldsymbol{0}, \boldsymbol{X})$. In particular we will be interested in drawing a standard Gaussian restricted to a subspace $H \subseteq \mathbb{R}^m$. The distribution of such a Gaussian is exactly $N(\mathbf{0}, \boldsymbol{X})$ where $\boldsymbol{X} = \sum_{i=1}^{\dim(H)} \boldsymbol{u}_i \boldsymbol{u}_i^\top$ and $\boldsymbol{u}_1, \dots, \boldsymbol{u}_{\dim(H)}$ is an orthonormal basis of H. The following properties are well known:

Lemma 113. *Let $H \subseteq \mathbb{R}^m$ be a subspace and let $N(\mathbf{0}, \boldsymbol{X})$ be the distribution of a standard Gaussian restricted to that subspace. Then for $\boldsymbol{y} \sim N(\mathbf{0}, \boldsymbol{X})$ one has (i) $\boldsymbol{y} \in H$ always; (ii) $\mathbb{E}[\|\boldsymbol{y}\|_2^2] = $ Tr$[\boldsymbol{X}] = \dim(H)$; (iii) $\mathbb{E}[y_i^2] \leq 1$ for all $i \in [m]$; (iv) Var$[\langle \boldsymbol{y}, \boldsymbol{b} \rangle] = \mathbb{E}[\langle \boldsymbol{y}, \boldsymbol{b} \rangle^2] \leq \|\boldsymbol{b}\|_2^2$ for all $\boldsymbol{b} \in \mathbb{R}^m$; (v) for any matrices $\boldsymbol{W}^1, \dots, \boldsymbol{W}^m \in \mathbb{R}^{n \times n}$ one has $\mathbb{E}[\| \sum_{i=1}^m y_i \boldsymbol{W}^i \|_F^2] \leq \sum_{i=1}^m \|\boldsymbol{W}^i\|_F^2$.*

The only property that is non-standard is (v). But note that we can use (iv) to justify that for each entry (k, ℓ) of the matrices one has $\mathbb{E}[(\sum_{i=1}^m y_i W_{k\ell}^i)^2] \leq \sum_{i=1}^m (W_{k\ell}^i)^2$; the claim then follows by linearity of expectation and summing over all entries $(k, \ell) \in [n]^2$.

Linear Algebra. For the analysis, we need an estimate on the trace of the product of symmetric matrices. The proof takes some care due to the non-commuting matrices. To get some intuition, consider the case when $\boldsymbol{A}_1, \boldsymbol{A}_2, \boldsymbol{B}$ are all diagonal matrices. In this case one can write $\boldsymbol{A}_1 \boldsymbol{B} = \text{diag}(\boldsymbol{a}_1)$ and $\boldsymbol{A}_2 \boldsymbol{B} = \text{diag}(\boldsymbol{a}_2)$ for some vectors $\boldsymbol{a}_1, \boldsymbol{a}_2 \in \mathbb{R}^n$ and the inequality simplifies to $\text{Tr}[\boldsymbol{A}_1 \boldsymbol{B} \boldsymbol{A}_2 \boldsymbol{B}] = \langle \boldsymbol{a}_1, \boldsymbol{a}_2 \rangle \leq \|\boldsymbol{a}\|_1 \cdot \|\boldsymbol{a}_2\|_1 = \text{Tr}[\boldsymbol{A}_1 |\boldsymbol{B}|] \cdot \text{Tr}[\boldsymbol{A}_2 |\boldsymbol{B}|]$ which is obviously true. Note that in the setting of [31] we would apply Lemma 114 with $\text{rank}(\boldsymbol{B}) = 2$, in which case the inequality can be tight up to constant factors. But in a different application with higher-rank matrices one could imagine a Cauchy-Schwarz or Hölder-type inequality yielding improved bounds.

Lemma 114. *Let $\boldsymbol{A}_1, \boldsymbol{A}_2, \boldsymbol{B} \in \mathbb{R}^{n \times n}$ be symmetric matrices with $\boldsymbol{A}_1, \boldsymbol{A}_2 \succeq 0$. Then*

$$\text{Tr}[\boldsymbol{A}_1 \boldsymbol{B} \boldsymbol{A}_2 \boldsymbol{B}] \leq \text{Tr}[\boldsymbol{A}_1 |\boldsymbol{B}|] \cdot \text{Tr}[\boldsymbol{A}_2 |\boldsymbol{B}|].$$

Proof. Write the spectral decomposition $\boldsymbol{B} = \sum_{i\in[n]} \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$. Then

$$\begin{aligned}
\mathrm{Tr}[\boldsymbol{A}_1\boldsymbol{B}\boldsymbol{A}_2\boldsymbol{B}] &= \sum_{i,j\in[n]} \lambda_i\lambda_j \cdot \big(\boldsymbol{v}_i^\top \boldsymbol{A}_1 \boldsymbol{v}_j\big)\big(\boldsymbol{v}_i^\top \boldsymbol{A}_2 \boldsymbol{v}_j\big) \\
&\leq \sum_{i,j\in[n]} |\lambda_i| \cdot |\lambda_j| \cdot \big\|\boldsymbol{A}_1^{1/2}\boldsymbol{v}_i\big\|_2 \cdot \big\|\boldsymbol{A}_2^{1/2}\boldsymbol{v}_i\big\|_2 \cdot \big\|\boldsymbol{A}_1^{1/2}\boldsymbol{v}_j\big\| \cdot \big\|\boldsymbol{A}_2^{1/2}\boldsymbol{v}_j\big\|_2 \\
&\leq \sum_{i,j\in[n]} |\lambda_i| \cdot |\lambda_j| \cdot \frac{1}{2}\Big(\|\boldsymbol{A}_1^{1/2}\boldsymbol{v}_i\|_2^2 \cdot \|\boldsymbol{A}_2^{1/2}\boldsymbol{v}_j\|^2 + \|\boldsymbol{A}_1^{1/2}\boldsymbol{v}_j\|_2^2 \cdot \|\boldsymbol{A}_2^{1/2}\boldsymbol{v}_i\|_2^2\Big) \\
&= \sum_{i,j\in[n]} |\lambda_i||\lambda_j| \cdot \frac{1}{2}\Big(\big(\boldsymbol{v}_i^\top \boldsymbol{A}_1 \boldsymbol{v}_i\big)\cdot\big(\boldsymbol{v}_j^\top \boldsymbol{A}_2 \boldsymbol{v}_j\big) + \big(\boldsymbol{v}_i^\top \boldsymbol{A}_2 \boldsymbol{v}_i\big)\cdot\big(\boldsymbol{v}_j^\top \boldsymbol{A}_1 \boldsymbol{v}_j\big)\Big) \\
&= \frac{1}{2}\Big(\mathrm{Tr}[\boldsymbol{A}_1|\boldsymbol{B}|]\cdot\mathrm{Tr}[\boldsymbol{A}_2|\boldsymbol{B}|] + \mathrm{Tr}[\boldsymbol{A}_2|\boldsymbol{B}|]\cdot\mathrm{Tr}[\boldsymbol{A}_1|\boldsymbol{B}|]\Big) \\
&= \mathrm{Tr}[\boldsymbol{A}_1|\boldsymbol{B}|]\cdot\mathrm{Tr}[\boldsymbol{A}_2|\boldsymbol{B}|],
\end{aligned}$$

where the first inequality is Cauchy-Schwarz and the second is AM-GM. $\square$

We also need a Taylor approximation for the trace of the inverse of a matrix. Again, it takes some care to handle the non-commutativity:

Lemma 115. *Let $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{S} \in \mathbb{R}^{n\times n}$ be symmetric matrices with $\boldsymbol{A}, \boldsymbol{S} \succ \boldsymbol{0}$ and $\|\delta \boldsymbol{A}^{-1}\boldsymbol{B}\|_{op} \leq \frac{1}{2}$. Then there is a value $c := c(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{S}, \delta) \in [0, 2]$ so that*

$$\mathrm{Tr}[(\boldsymbol{A} - \delta\boldsymbol{B})^{-1}\boldsymbol{S}] = \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{S}] + \delta\mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}] + c\delta^2\mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}].$$

Proof. We abbreviate $\boldsymbol{M} := \delta\boldsymbol{A}^{-1}\boldsymbol{B}$. As $\|\boldsymbol{M}\|_{\mathrm{op}} \leq \frac{1}{2}$, the matrix $\boldsymbol{I}_n - \boldsymbol{M}$ is non-singular. We obtain

$$(\boldsymbol{A} - \delta\boldsymbol{B})^{-1} = (\boldsymbol{A}(\boldsymbol{I}_n - \boldsymbol{M}))^{-1} = (\boldsymbol{I}_n - \boldsymbol{M})^{-1}\boldsymbol{A}^{-1} = \sum_{k=0}^{\infty} \boldsymbol{M}^k\boldsymbol{A}^{-1},$$

so that

$$\mathrm{Tr}[(\boldsymbol{A} - \delta\boldsymbol{B})^{-1}\boldsymbol{S}] = \sum_{k=0}^{\infty}\mathrm{Tr}[\boldsymbol{M}^k\boldsymbol{A}^{-1}\boldsymbol{S}] = \sum_{k=0}^{2}\mathrm{Tr}[\boldsymbol{M}^k\boldsymbol{A}^{-1}\boldsymbol{S}] + \sum_{k=3}^{\infty}\mathrm{Tr}[\boldsymbol{M}^k\boldsymbol{A}^{-1}\boldsymbol{S}].$$

For any $k \geq 3$,

$$
\begin{aligned}
|\mathrm{Tr}[\boldsymbol{M}^k \boldsymbol{A}^{-1}\boldsymbol{S}]| &= \delta^k \cdot \left|\mathrm{Tr}\left[(\boldsymbol{A}^{-1/2}\boldsymbol{B}\boldsymbol{A}^{-1/2})^{k-2} \cdot \boldsymbol{A}^{-1/2}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1/2}\right]\right| \\
&\leq \delta^k \|(\boldsymbol{A}^{-1/2}\boldsymbol{B}\boldsymbol{A}^{-1/2})^{k-2}\|_{\mathrm{op}} \cdot \mathrm{Tr}[\boldsymbol{A}^{-1/2}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1/2}\Big] \\
&\leq \|\boldsymbol{M}\|_{\mathrm{op}}^{k-2} \cdot \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}].
\end{aligned}
$$

We conclude

$$
\Big|\sum_{k=3}^{\infty} \mathrm{Tr}[\boldsymbol{M}^k \boldsymbol{A}^{-1}\boldsymbol{S}]\Big| \leq \sum_{j=1}^{\infty} 2^{-j} \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}] = \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{S}],
$$

so that the statement follows. □

6.3 *Main technical result*

We now show our main result, Theorem 105. Fix symmetric matrices $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m \in \mathbb{R}^{n \times n}$ with $\boldsymbol{S} := \sum_{i=1}^m |\boldsymbol{A}_i|$ satisfying $\boldsymbol{S} \preceq \boldsymbol{I}_n$ and $\mathrm{Tr}[\boldsymbol{S}] = \tau \geq 1$ and set $\varepsilon > 0$ so that $m = \frac{\tau}{\varepsilon^2}$. Let K be the body as defined in Theorem 105 and fix a parameter $\alpha > 0$. Ideally, the goal would be to prove that a random Gaussian from $N(\boldsymbol{0}, \boldsymbol{I}_m)$ is on average close to K. Instead, we prove that there is a random variable $\boldsymbol{x}$ that is close to a Gaussian and ends up in K with high probability. The strategy is to generate such a near-Gaussian random variable $\boldsymbol{x}$ by performing a *Brownian motion* that adds up independent Gaussians $\boldsymbol{y}^{(t)}$ with a tiny step size δ. The key ingredient is that in each iteration t we walk inside a subspace of dimension at least $(1-\alpha^2)m$, meaning that we draw $\boldsymbol{y}^{(t)} \sim N(\boldsymbol{0}, \boldsymbol{X}^{(t)})$ with tr$[\boldsymbol{X}^{(t)}] \geq (1-\alpha^2)m$. This can be understood as blocking the movement in $\alpha^2 m$ dimensions that are "dangerous". Then the expected Euclidean distance of the outcome $\boldsymbol{x} = \delta \sum_{t=1}^{1/\delta^2} \boldsymbol{y}^{(t)}$ to an unrestricted Gaussian is at most $\alpha\sqrt{m}$. It remains to argue that the subspace can be chosen so that at the end of the Brownian motion, $\boldsymbol{x}$ ends up in K. For this sake we define a *potential function*

$$
\boldsymbol{A}_{C,D}(\boldsymbol{x}) := (C + D\|\boldsymbol{x}\|_2^2) \cdot \boldsymbol{I}_n - \sum_{i=1}^m x_i \boldsymbol{A}_i \quad \text{and} \quad \Phi_{C,D}(\boldsymbol{x}) := \frac{1}{2}\mathrm{Tr}\left[\left(\boldsymbol{S} + \frac{\tau}{n} \cdot \boldsymbol{I}_n\right) \cdot \boldsymbol{A}_{C,D}(\boldsymbol{x})^{-1}\right]
$$

We initialize the random walk with $\boldsymbol{x} := \boldsymbol{0}$ so that $\boldsymbol{A}_{C,D}(\boldsymbol{x}) \succ 0$. If the update steps are small and we keep the potential function $\Phi_{C,D}(\boldsymbol{x})$ bounded, we can infer that $\sum_{i=1}^m x_i \boldsymbol{A}_i \preceq$

$(C + D\|\boldsymbol{x}\|_2^2) \cdot \boldsymbol{I}_n$ at any given time. More precisely we show that, for a particular choice of parameters $C, D > 0$ (later we will choose $C = \Theta(\frac{\varepsilon}{\alpha})$ and $D = \Theta(\frac{\varepsilon}{\alpha m})$), an update of $\boldsymbol{x}' = \boldsymbol{x} + \delta \boldsymbol{y}^{(t)}$ in expectation does not increase the value of the potential function — assuming that the current value of the potential function is small enough and $\boldsymbol{y}^{(t)}$ is taken from the aforementioned subspace.

There is the technical issue that the potential function goes up to ∞ as the minimal eigenvalue of $\boldsymbol{A}_{C,D}(\boldsymbol{x})$ approaches 0. We solve this problem by defining another distribution $N_{\leq m}(\boldsymbol{0}, \boldsymbol{X})$ that draws $\boldsymbol{y} \sim N(\boldsymbol{0}, \boldsymbol{X})$, but if $\|\boldsymbol{y}\|_2 > m$, then $\boldsymbol{y}$ is replaced with $\boldsymbol{0}$. Recall that by Corollary 5 one has $\Pr_{\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{X})}[\|\boldsymbol{y}\|_2 > m] \leq 2^{-m}$ for any $\boldsymbol{X} \preceq \boldsymbol{I}_m$. A second problem is that keeping the potential function low in expectation is not sufficient — if the potential function ever crosses a certain threshold, the analysis stops working. However, a single step in the Brownian motion can be analyzed as follows:

Lemma 116. *Fix $0 < \alpha < 1$ and $m \geq \max\left\{200, 9\log n, \frac{49}{\alpha^2}\right\}$. Let $\boldsymbol{x} \in \mathbb{R}^m$ and suppose $\boldsymbol{A}_{C,D}(\boldsymbol{x}) \succ 0$, $\Phi_{C,D}(\boldsymbol{x}) \leq \frac{Dm^2\alpha^2}{10}$ as well as $\delta := \frac{\alpha}{45m^3n}$. Define $F(\boldsymbol{y})$ as the unique value for which*

$$\Phi_{C+\delta^2 F(\boldsymbol{y}),D}(\boldsymbol{x} + \delta\boldsymbol{y}) = \Phi_{C,D}(\boldsymbol{x}).$$

Then there is a covariance matrix $\boldsymbol{X} \in \mathbb{R}^{m\times m}$ with $0 \preceq \boldsymbol{X} \preceq \boldsymbol{I}_m$ and Tr$[\boldsymbol{X}] \geq (1-\alpha^2)\cdot m$ so that $\mathbb{E}_{\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}[F(\boldsymbol{y})] \leq 0$ while always $|F(\boldsymbol{y})| \leq 4Dm^4$. Further, $\boldsymbol{A}_{C+\delta^2F(\boldsymbol{y}),D}(\boldsymbol{x}+\delta\boldsymbol{y}) \succ 0$.

We postpone the proof of this lemma to Section 4. First, we show how we can use it to obtain the main theorem:

Proof of Theorem 105. Let $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m \in \mathbb{R}^{n\times n}$ be symmetric matrices with $\sum_{i=1}^m |\boldsymbol{A}_i| \preceq \boldsymbol{I}_n$ so that $m = \frac{\tau}{\varepsilon^2} \geq 200$. Fix a parameter $0 < \alpha < 1$ and keep in mind that the goal is to prove that $\gamma_m\left(\frac{50}{\alpha}K + \alpha\sqrt{m}B_2^m\right) \geq \frac{1}{2}$. First we justify that we may assume $\tau \geq 1$. If the result holds for $\tau = 1$, then for any $\tau < 1$ we may consider instead the matrices $\boldsymbol{A}_i' := \boldsymbol{A}_i/\tau$ which still satisfy $\sum_{i=1}^m |\boldsymbol{A}_i'| \preceq \boldsymbol{I}_n$ and $\sum_{i=1}^n \mathrm{Tr}[|\boldsymbol{A}_i'|] = 1$. Then in order to show a measure lower bound for the discrepancy body associated with $\boldsymbol{A}_i$, it suffices to do so for the (smaller) body associated with $\boldsymbol{A}_i'$. Hence we may assume $\tau \geq 1$ from now on. Assume at first $m \geq \frac{49}{\alpha^2}$.

Note that the potential function $\Phi_{C,D}(\boldsymbol{x})$ is *one-sided* in the sense that it only controls the *maximum* eigenvalue of $\sum_{i=1}^m x_i \boldsymbol{A}_i$. For this sake we abbreviate

$$\tilde{\boldsymbol{A}}_i := \begin{pmatrix} \boldsymbol{A}_i & \boldsymbol{0} \\ \boldsymbol{0} & -\boldsymbol{A}_i \end{pmatrix} \in \mathbb{R}^{2n \times 2n}$$

Note that this allows us to rewrite $K = \left\{ \boldsymbol{x} \in \mathbb{R}^m \mid \sum_{i=1}^m x_i \tilde{\boldsymbol{A}}_i \preceq \varepsilon \boldsymbol{I}_{2n} \right\}$. In wise foresight we choose $D := \frac{2\varepsilon}{\alpha m}$ and define C so that $\frac{2\tau}{C} = \frac{Dm^2\alpha^2}{10}$ which results in $C = \frac{10\varepsilon}{\alpha}$. We define a small enough step size of $\delta := \frac{\alpha}{45m^3 n}$ and choose $T := \frac{1}{\delta^2}$ as the number of iterations.

Note that by definition $\Phi_{C,D}(\boldsymbol{0}) = \frac{2\tau}{C} = \frac{Dm^2\alpha^2}{10}$. Consider the following (hypothetical) algorithm:

(1) Set $\boldsymbol{x}^{(0)} := \boldsymbol{0}$

(2) For $t = 1$ TO T do

  (3) Apply Lemma 116 for $\boldsymbol{x}^{(t-1)}$ and let $\boldsymbol{X}^{(t)}$ be the obtained covariance matrix.

  (4) Sample $\boldsymbol{y}^{(t)} \sim N(\boldsymbol{0}, \boldsymbol{X}^{(t)})$ and $\boldsymbol{z}^{(t)} \sim N(\boldsymbol{0}, \boldsymbol{I}_m - \boldsymbol{X}^{(t)})$.
  If $\|\boldsymbol{y}^{(t)}\|_2 \leq m$ then $(\boldsymbol{y}^{(t)}_{\leq m}, \boldsymbol{y}^{(t)}_{>m}) := (\boldsymbol{y}^{(t)}, \boldsymbol{0})$, otherwise $(\boldsymbol{y}^{(t)}_{\leq m}, \boldsymbol{y}^{(t)}_{>m}) := (\boldsymbol{0}, \boldsymbol{y}^{(t)})$.

  (5) Update $\boldsymbol{x}^{(t)} := \boldsymbol{x}^{(t-1)} + \delta \boldsymbol{y}^{(t)}_{\leq m}$.

At the end, let $\boldsymbol{Y} := \boldsymbol{x}^{(T)} = \delta \sum_{t=1}^T \boldsymbol{y}^{(t)}$ and $\boldsymbol{Z} := \delta \sum_{t=1}^T \boldsymbol{z}^{(t)}$. Note that $\boldsymbol{Y} + \boldsymbol{Z} \sim N(\boldsymbol{0}, \boldsymbol{I}_m)$.

Claim. *The following events all hold simultaneously with probability at least $\frac{1}{2}$:*

(a) *One has $\boldsymbol{y}^{(t)}_{>m} = \boldsymbol{0}$ for all $t = 1, \ldots, T$*

(b) *One has $\delta^2 \sum_{t=1}^T F(\boldsymbol{y}^{(t)}_{\leq m}) \leq \frac{C}{10}$*

(c) *One has $\|\boldsymbol{Z}\|_2^2 \leq 5\alpha^2 m$*

(d) *One has $\|\boldsymbol{Y}\|_2^2 \leq 5m$*

Proof of claim. By Corollary 5, the failure probability for (a) is bounded by

$$T \cdot 2^{-m} = \frac{1}{\delta^2} \cdot 2^{-m} \leq \frac{1}{25},$$

which follows as

$$\delta = \frac{\alpha}{45m^3n^2} \geq \frac{\sqrt{49/m}}{45m^3 \cdot 2^{m/3}} \overset{m\geq 200}{\geq} 5 \cdot 2^{-m/2}.$$

For (b), note that for every step t, the conditional expectation of $F(\boldsymbol{y}_{\leq m}^{(t)})$ is nonpositive, and $|F(\boldsymbol{y}_{\leq m}^{(t)})| \leq 7Dm^4n$. Then using Azuma's inequality and $\frac{C}{D} = 5m$, one has

$$\Pr\left[\delta^2 \sum_{t=1}^{T} F(\boldsymbol{y}_{\leq m}^{(t)}) > \frac{C}{10}\right] \leq \exp\left(-\frac{1}{2}\cdot\frac{(\frac{C}{10}\delta)^2}{(\delta^2 \cdot 7Dm^4n)^2}\right) = \exp\left(-\frac{1}{392\delta^2m^6n^2}\right) \leq \exp(-3) \leq \frac{1}{20},$$

since $\delta^2 \leq \frac{1}{1176m^6n^2}$.

For (c), note that $\mathbb{E}[\|\boldsymbol{Z}\|_2^2] = \sum_{t=1}^{T} \mathrm{Tr}[\delta^2 \cdot (\boldsymbol{I}_m - \boldsymbol{X}^{(t)})] \leq \alpha^2 m$, so by Markov's inequality $\Pr[\|\boldsymbol{Z}\|_2^2 > 5\alpha^2 m] \leq \frac{1}{5}$. Similarly, $\mathbb{E}[\|\boldsymbol{Y}\|_2^2] \leq m$, so $\|\boldsymbol{Y}\|_2^2 > 5m$ with probability at most $\frac{1}{5}$. The total failure probability is therefore at most $\frac{1}{25} + \frac{1}{20} + \frac{1}{5} + \frac{1}{5} = 0.49 < \frac{1}{2}$. $\diamond$

If the events in the claim hold, we have $\|\boldsymbol{Z}\|_2 \leq \alpha\sqrt{5m}$, $\boldsymbol{A}_{1.1C,D}(\boldsymbol{Y}) \succ 0$ and

$$\sum_{i=1}^{m} Y_i \boldsymbol{A}_i \preceq (1.1C + D\|\boldsymbol{Y}\|_2^2) \cdot \boldsymbol{I}_{2n} \preceq \left(1.1 \cdot \frac{10\varepsilon}{\alpha} + \frac{10\varepsilon}{\alpha}\right) = \frac{21\varepsilon}{\alpha} \cdot \boldsymbol{I}_{2n}.$$

It remains to finish the arguments behind the proof strategy. By a slight abuse of notation, let $\gamma_m(\boldsymbol{x}) := \frac{1}{(2\pi)^{m/2}} e^{-\|\boldsymbol{x}\|_2^2/2}$ be the density of the Gaussian at a point $\boldsymbol{x}$. We define $p(\boldsymbol{x})$ as the conditional probability that the properties (a)-(d) are satisfied, conditioned on the event that $\boldsymbol{Y} + \boldsymbol{Z} = \boldsymbol{x}$. Then our reasoning above has proven that $\int_{\mathbb{R}^m} \gamma_m(\boldsymbol{x}) \cdot p(\boldsymbol{x}) d\boldsymbol{x} \geq \frac{1}{2}$. Now define the set $Q := \{\boldsymbol{x} \in \mathbb{R}^m \mid p(\boldsymbol{x}) > 0\}$. As $0 \leq p(\boldsymbol{x}) \leq 1$ we must have $\gamma_m(Q) \geq \frac{1}{2}$. By construction, for every $\boldsymbol{x} \in Q$, there is at least one witness outcome $\boldsymbol{Y} + \boldsymbol{Z} = \boldsymbol{x}$ so that $\boldsymbol{Y} \in \frac{21\varepsilon}{\alpha} K$ and $\|\boldsymbol{Z}\|_2 \leq \alpha\sqrt{5m}$. Then a slight reparametrization of $\alpha' := \sqrt{5}\alpha$ gives the claim as $21\sqrt{5} < 50$.

One final detail is that in Lemma 116 we assume $m \geq \frac{49}{\alpha^2}$. Assume now that $\alpha \leq \frac{7}{\sqrt{m}}$. As $\mathbb{E}_{\boldsymbol{x}\sim N(\boldsymbol{0},\boldsymbol{I}_m)} \|\sum_{i=1}^{m} x_i \boldsymbol{A}_i\|_{\mathrm{op}}^2 \leq \mathbb{E}_{\boldsymbol{x}\sim N(\boldsymbol{0},\boldsymbol{I}_m)} \|\sum_{i=1}^{m} x_i \boldsymbol{A}_i\|_F^2 = \sum_{i=1}^{n} \mathrm{Tr}[\boldsymbol{A}_i^2] \leq \sum_{i=1}^{n} \mathrm{Tr}[\boldsymbol{A}_i] = \tau$, by Markov's inequality it follows that $\|\sum_{i=1}^{m} x_i \boldsymbol{A}_i\|_{\mathrm{op}}^2 \leq 2\tau$ with probability at least $1/2$. Thus we get a measure lower bound $\gamma_m(\sqrt{2m}K) \geq 1/2$.

In particular, for $\alpha \leq \frac{7}{\sqrt{m}}$, $\gamma_m(\frac{50}{\alpha}K + \alpha\sqrt{m}B_2^m) \geq \gamma_m(\frac{50}{\alpha}K) \geq \gamma_m(\sqrt{2m}K) \geq \frac{1}{2}$. $\square$

6.4 Analysis of a single step

In this section, we prove Lemma 116 and some variants that will be needed later.

Proof of Lemma 116. To simplify notation, we abbreviate matrices

$$\boldsymbol{A} := \boldsymbol{A}_{C,D}(\boldsymbol{x}), \quad \tilde{\boldsymbol{B}} := \sum_{i=1}^{m} y_i \boldsymbol{A}_i, \boldsymbol{B} := \tilde{\boldsymbol{B}} - \delta(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))\boldsymbol{I}_n,$$

$$\boldsymbol{S} := \sum_{i=1}^{m} |\boldsymbol{A}_i| \quad \text{and} \quad \tilde{\boldsymbol{S}} := \frac{1}{2}\Big(\boldsymbol{S} + \frac{\tau}{n} \cdot \boldsymbol{I}_n\Big).$$

Next, we define an index set

$$\mathcal{I} := \Big\{ i \in [m] : \mathrm{Tr}\big[\boldsymbol{A}^{-1}|\boldsymbol{A}_i|\big] \leq \frac{1.1}{\alpha^2 m} \cdot \mathrm{Tr}\big[\boldsymbol{S}\boldsymbol{A}^{-1}\big] \Big\}$$

Here $[m] \setminus \mathcal{I}$ are the "dangerous" indices in the sense that updating $\boldsymbol{x}$ in these coordinates might disproportionally change the potential function. Note that by Markov's inequality, we have $|\mathcal{I}| \geq (1 - \frac{\alpha^2}{1.1})m$. Consider the subspace

$$\begin{aligned} H := \Big\{ \boldsymbol{y} \in \mathbb{R}^m : y_i = 0 \ \forall i \notin \mathcal{I}, \langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0, & \sum_{i=1}^{m} y_i \cdot \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] = 0, \\ & \sum_{i=1}^{m} y_i \cdot \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}] = 0, \\ & \sum_{i=1}^{m} y_i \cdot \mathrm{Tr}[\boldsymbol{A}^{-2}\boldsymbol{A}_i\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] = 0 \Big\} \end{aligned}$$

so that $\dim(H) \geq |\mathcal{I}| - 4 \geq (1-\alpha^2)m$ for $m \geq \frac{49}{\alpha^2}$. Further, $\dim(H) \geq |\mathcal{I}| - 4 \geq 0.47m$ for $m \geq 100$. We choose $\boldsymbol{X}$ so that $N(\boldsymbol{0}, \boldsymbol{X})$ is the standard Gaussian restricted to H.

The remaining proof is organized in 4 claims, where Claim I-III justify that the Taylor approximation is well behaved while Claim IV contains a very crucial upper bound. We begin by showing a rather crude upper bound on $|F(\boldsymbol{y})|$ for $\|\boldsymbol{y}\|_2 \leq m$.

Claim I. *For $\boldsymbol{y}$ with $\|\boldsymbol{y}\|_2 \leq m$, one has $|F(\boldsymbol{y})| \leq \frac{2m}{\delta}$, $\|\boldsymbol{B}\|_{op} \leq 4m$ and $\|\delta \boldsymbol{A}^{-1}\boldsymbol{B}\|_{op} < \frac{1}{2}$.*

Proof of Claim I. Note that in order for the potential functions $\Phi_{C+\delta^2 F(\boldsymbol{y}),D}(\boldsymbol{x} + \delta\boldsymbol{y})$ and

$\Phi_{C,D}(\boldsymbol{x})$ to be identical, we know that the difference matrix

$$\begin{aligned}\boldsymbol{A}_{C+\delta^2 F(\boldsymbol{y}),D}(\boldsymbol{x}+\delta\boldsymbol{y}) - \boldsymbol{A}_{C,D}(\boldsymbol{x}) &= \delta^2(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))\cdot \boldsymbol{I}_n - \delta\sum_{i=1}^m y_i\boldsymbol{A}_i \\ &\preceq \delta^2(Dm^2 + F(\boldsymbol{y}))\cdot\boldsymbol{I}_n + \delta\|\boldsymbol{y}\|_\infty \underbrace{\sum_{i=1}^m |\boldsymbol{A}_i|}_{\preceq \boldsymbol{I}_n} \\ &\preceq \big(\delta Dm^2 + \delta F(\boldsymbol{y}) + m\big)\cdot\delta\cdot\boldsymbol{I}_n\end{aligned}$$

must have one eigenvalue at least 0 and one eigenvalue at most 0. There would be no positive eigenvalues if $\delta F(\boldsymbol{y}) < -2m < -\delta Dm^2 - m$, and similarly no negative eigenvalues if $\delta F(\boldsymbol{y}) > 2m$. Hence we conclude $|F(\boldsymbol{y})| \leq \frac{2m}{\delta}$. This bound is good enough to show that

$$\|\boldsymbol{B}\|_{\text{op}} \leq \underbrace{\Big\|\sum_{i=1}^m y_i\boldsymbol{A}_i\Big\|_{\text{op}}}_{\leq m} + \delta\Big(D\|\boldsymbol{y}\|_2^2 + \frac{2m}{\delta}\Big) \leq 3m + \delta Dm^2 < 4m.$$

Since $\|\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\|_{\text{op}} \leq \Phi_{C,D}(\boldsymbol{x}) \leq \frac{Dm^2}{10}$ and $\|\tilde{\boldsymbol{S}}^{-1}\|_{\text{op}} \leq \frac{2n}{\tau} \leq 2n$, it follows we may also bound $\|\boldsymbol{A}^{-1}\|_{\text{op}} \leq \|\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\|_{\text{op}}\cdot\|\tilde{\boldsymbol{S}}^{-1}\|_{\text{op}} \leq \frac{Dm^2 n}{5}$, so that $\|\delta\boldsymbol{A}^{-1}\boldsymbol{B}\|_{\text{op}} \leq \delta\|\boldsymbol{A}^{-1}\|_{\text{op}}\|\boldsymbol{B}\|_{\text{op}} < \frac{1}{2}$ since $\delta = \frac{\alpha}{45m^3n^2} < \frac{1}{2}\cdot\frac{1}{4m}\cdot\frac{5}{Dm^2n}$ as $D = \frac{2}{\alpha m}\sqrt{\frac{\tau}{m}} < \frac{n}{\alpha}$. ◇

Now we can apply the matrix Taylor approximation from Lemma 115 and use that for every $\boldsymbol{y} \in H$ with $\|\boldsymbol{y}\|_2 \leq m$, there exists some $|c| \leq 2$ such that the difference in the potential function is

$$\begin{aligned}0 &\overset{\text{Def } F(\boldsymbol{y})}{=} \Phi_{C+\delta^2 F(\boldsymbol{y}),D}(\boldsymbol{x}+\delta\boldsymbol{y}) - \Phi_{C,D}(\boldsymbol{x}) \\ &= \operatorname{Tr}\Big[(\boldsymbol{A}-\delta\boldsymbol{B})^{-1}\tilde{\boldsymbol{S}}\Big] - \operatorname{Tr}\Big[\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}\Big] \\ &\overset{\text{Lem 115}}{=} \delta\cdot\operatorname{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] + c\delta^2\cdot\operatorname{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] \\ &= -\delta^2(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))\cdot\operatorname{Tr}[\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}] + \delta\underbrace{\sum_{i=1}^m y_i\operatorname{Tr}[\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]}_{=0} + c\delta^2\cdot\operatorname{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] \\ &\overset{\boldsymbol{y}\in H}{=} \delta^2\Big(-(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))\cdot\operatorname{Tr}[\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}] + c\cdot\operatorname{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]\Big). \qquad (**)\end{aligned}$$

Observe that in the last equation we have conveniently used that due to the linear constraints defining H, we have $\operatorname{Tr}[\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] = 0$ for all $\boldsymbol{y} \in H$. Now we can show that the quantity $F(\boldsymbol{y})$ is a lot smaller than we have proven so far — in fact its maximum length is independent of the step size δ:

Claim II. *For every $\boldsymbol{y} \in H$ with $\|\boldsymbol{y}\|_2 \leq m$ one has $|F(\boldsymbol{y})| \leq 7Dm^4n$.*

Proof of Claim II. We rearrange $(**)$ for $F(\boldsymbol{y})$ and obtain

$$|F(\boldsymbol{y})| \leq \frac{|c| \cdot \mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]}{\mathrm{Tr}[\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}]} + D\underbrace{\|\boldsymbol{y}\|_2^2}_{\leq m^2} \leq 2 \cdot \|\boldsymbol{A}^{-1}\|_{\mathrm{op}} \cdot \|\boldsymbol{B}\|_{\mathrm{op}}^2 + Dm^2 \leq 7Dm^4n,$$

using the estimates $\|\boldsymbol{B}\|_{\mathrm{op}} \leq 4m$ and $\|\boldsymbol{A}^{-1}\|_{\mathrm{op}} \leq \|\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\|_{\mathrm{op}} \cdot \|\tilde{\boldsymbol{S}}^{-1}\|_{\mathrm{op}} \leq \frac{Dm^2n}{5\tau} \leq \frac{Dm^2n}{5}$. $\diamond$

Next, we justify that $\mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] \approx \mathrm{Tr}[\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]$ up to lower order terms.

Claim III. *For any $\boldsymbol{y} \in H$ with $\|\boldsymbol{y}\|_2 \leq m$ one has*

$$\Big|\mathrm{Tr}[\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}] - \mathrm{Tr}[\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]\Big| \leq \delta^2 \cdot 12D^3m^{10}n^3 \cdot \mathrm{Tr}[\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}].$$

Proof of Claim III. Since $\boldsymbol{B} = \tilde{\boldsymbol{B}} - \lambda\boldsymbol{I}_n$ for $\lambda := \delta(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))$, the difference equals

$$\Big| - \lambda(\underbrace{\mathrm{Tr}[\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-2}\tilde{\boldsymbol{S}}]}_{=0} + \underbrace{\mathrm{Tr}[\boldsymbol{A}^{-2}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}]}_{=0}) + \lambda^2\mathrm{Tr}[\boldsymbol{A}^{-3}\tilde{\boldsymbol{S}}]\Big| \leq \delta^2\mathrm{Tr}[\boldsymbol{A}^{-3}\tilde{\boldsymbol{S}}] \cdot 60(Dm^4n)^2.$$

Here we use $\lambda^2 = \delta^2(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y}))^2 \leq \delta^2(7.7Dm^4n)^2$. We have also made use of the linear constraints in the choice of the subspace H. It remains to use that $\mathrm{Tr}[\boldsymbol{A}^{-3}\tilde{\boldsymbol{S}}] \leq \|\boldsymbol{A}^{-1}\|_{\mathrm{op}} \cdot \mathrm{Tr}[\boldsymbol{A}^{-2}\boldsymbol{S}]$ with the bound $\|\boldsymbol{A}^{-1}\|_{\mathrm{op}} \leq \frac{Dm^2n}{5}$. $\diamond$

Now we prove the central core of this theorem: in expectation for a Gaussian $\boldsymbol{y}$ from the subspace H, the quadratic term $\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}]$ is bounded by a term that we can offset in the potential function by the length increase of $\boldsymbol{x}$.

Claim IV. *One has $\mathbb{E}_{\boldsymbol{y} \sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}\left[\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}]\right] \leq \frac{1.1}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}] \cdot \mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-2}]$.*

Proof of Claim IV. The argument for this claim needs some care, as we have in general $\mathbb{E}[y_iy_j] \neq 0$ since we draw $\boldsymbol{y}$ from a subspace H. We abbreviate $\boldsymbol{W}_i := \tilde{\boldsymbol{S}}^{1/2}\boldsymbol{A}^{-1/2}\boldsymbol{A}_i\boldsymbol{A}^{-1} \in$

$\mathbb{R}^{n\times n}$ (note that these matrices will in general not be symmetric). Then

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}\Big[\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}\tilde{\boldsymbol{B}}\boldsymbol{A}^{-1}]\Big] &\overset{(i)}{=} \mathop{\mathbb{E}}_{\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}\Big[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} y_i y_j \mathrm{Tr}\Big[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-1}\boldsymbol{A}_j\boldsymbol{A}^{-1}\Big]\Big] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}\Big[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} y_i y_j \left\langle \boldsymbol{W}_i, \boldsymbol{W}_j\right\rangle_F\Big] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})}\Big[\Big\|\sum_{i\in\mathcal{I}} y_i\boldsymbol{W}_i\Big\|_F^2\Big] \\
&\overset{(ii)}{\leq} \sum_{i\in\mathcal{I}}\|\boldsymbol{W}_i\|_F^2 = \sum_{i\in\mathcal{I}}\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-1}\boldsymbol{A}_i\boldsymbol{A}^{-1}] \\
&\overset{\text{Lem 114}}{\leq} \sum_{i\in\mathcal{I}}\mathrm{Tr}\big[\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}|\boldsymbol{A}_i|\big]\cdot\mathrm{Tr}\big[\boldsymbol{A}^{-1}|\boldsymbol{A}_i|\big] \\
&\overset{(iii)}{\leq} \frac{1.1}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}]\sum_{i\in\mathcal{I}}\mathrm{Tr}\big[\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}|\boldsymbol{A}_i|\big] \\
&= \frac{1.1}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}]\mathrm{Tr}\Big[\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\cdot\underbrace{\sum_{i\in\mathcal{I}}|\boldsymbol{A}_i|}_{\preceq\boldsymbol{I}_n}\Big] \\
&\overset{\boldsymbol{A}^{-1}\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\succ 0}{\leq} \frac{1.1}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}]\cdot\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-2}].
\end{aligned}$$

In (i), we use that $y_i = 0$ for $i \notin \mathcal{I}$. In (ii) we use Lemma 113 with the subtlety that replacing $\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{X})$ by the capped sample $\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})$ can only decrease the length $\|\sum_{i\in\mathcal{I}} y_i\boldsymbol{W}_i\|_F^2$. In (iii) we use that we have selected the indices $\mathcal{I}$ so that $\mathrm{Tr}[\boldsymbol{A}^{-1}|\boldsymbol{A}_i|] \leq \frac{1.1}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}]$ for any $i\in\mathcal{I}$. $\diamond$

Now we have everything to finish the analysis. Taking expectation over $\boldsymbol{y}\sim N_{\leq m}(\boldsymbol{0},\boldsymbol{X})$ on both sides of $(**)$ gives

$$\begin{aligned}
0 &\overset{(**)}{=} -(D\,\mathbb{E}[\|\boldsymbol{y}\|_2^2] + \mathbb{E}[F(\boldsymbol{y})])\cdot\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-2}] + c\cdot\mathbb{E}[\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}\boldsymbol{A}^{-1}]] \\
&\overset{\text{Claims III \& IV}}{\leq} \Big(-0.49Dm - \mathbb{E}[F(\boldsymbol{y})] + \frac{2.2}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}] + \delta^2\cdot 24D^3m^{10}n^3\Big)\cdot\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-2}] \\
&\leq \Big(-0.44Dm - \mathbb{E}[F(\boldsymbol{y})] + \frac{2.2}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}]\Big)\cdot\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-2}]
\end{aligned}$$

In the first inequality we have used Claim III, and in the last inequality we used the fact that $\delta^2 = \frac{\alpha^2}{2025m^6n^4} \leq 0.05Dm\cdot\frac{1}{24D^3m^{10}n^3} = \frac{1}{480D^2m^9n^3} = \frac{\alpha^2}{1920m^6n^3}$. Here we also use that by

Corollary 5 one has $\mathbb{E}[\|\boldsymbol{y}\|_2^2] \geq \dim(H) \cdot \left(1 - 2^{-\dim(H)}\right) \geq 0.49m$ as $\dim(H) \geq 0.47m$ and $m \geq 200$. Combining the two above inequalities, we conclude

$$\mathbb{E}[F(\boldsymbol{y})] \leq \frac{2.2}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}] - 0.44Dm \leq \frac{2.2}{\alpha^2 m}\mathrm{Tr}[\boldsymbol{S}\boldsymbol{A}^{-1}] - 0.44Dm \leq \frac{4.4}{\alpha^2 m}\mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}] - 0.44Dm \leq 0,$$

making use of $\boldsymbol{S} \preceq 2\tilde{\boldsymbol{S}}$ and of the assumed bound on $\Phi_{C,D}(\boldsymbol{x}) = \mathrm{Tr}[\tilde{\boldsymbol{S}}\boldsymbol{A}^{-1}] \leq \frac{Dm^2\alpha^2}{10}$.

It remains to argue $\tilde{\boldsymbol{A}} := \boldsymbol{A}_{C+\delta^2 F(\boldsymbol{y}),D} \succ 0$. Recall from the proof of Claim I we have

$$\tilde{\boldsymbol{A}} - \boldsymbol{A} = \delta^2(D\|\boldsymbol{y}\|_2^2 + F(\boldsymbol{y})) \cdot \boldsymbol{I}_n - \delta\sum_{i=1}^m y_i\boldsymbol{A}_i \succeq (\delta F(\boldsymbol{y}) - m) \cdot \delta \cdot \boldsymbol{I}_n \succeq -3m \cdot \delta \cdot \boldsymbol{I}_n,$$

where we have used $\delta F(\boldsymbol{y}) \geq -2m$. Recall that the least eigenvalue of $\boldsymbol{A}$ is at least $\frac{5}{Dm^2 n}$. It follows the least eigenvalue of $\tilde{\boldsymbol{A}}$ is at least $\frac{5}{Dm^2 n} - \delta \cdot 3m > 0$ as $\delta = \frac{\alpha}{45m^3n^2} < \frac{1}{0.6Dm^3}$. $\square$

Finally, we can prove Theorem 107 and give an analysis of the full algorithm from Section 6.1.1. The basic intuition is that we start with a weight vector $\boldsymbol{s} := (1, \ldots, 1)$ so that $\sum_{i=1}^m s_i\boldsymbol{A}_i = \boldsymbol{I}_n$. Then in each iteration we find a partial coloring and use it to update the weights so that at least a constant fraction of the weights drop to 0.

Proof of Theorem 107. Consider one iteration of the algorithm where the current weights are $\boldsymbol{s} \in \mathbb{R}^m_{\geq 0}$. The body defined in step (3) is $K := \{\boldsymbol{x} \in \mathbb{R}^{\mathrm{supp}(\boldsymbol{s})} \mid \|\sum_{i\in\mathrm{supp}(\boldsymbol{s})} x_i s_i \boldsymbol{A}_i\|_{\mathrm{op}} \leq 1000\tilde{\varepsilon}\}$. Hence, by Theorem 106 and Theorem 7, we have a point $\boldsymbol{x}^* \in [-1,1]^{\mathrm{supp}(\boldsymbol{s})}$ with $\|\sum_{i=1}^m x_i^* s_i\boldsymbol{A}_i\|_{\mathrm{op}} \leq \Omega(\sqrt{\frac{n}{|\mathrm{supp}(\boldsymbol{s})|}})$ and at least $\Omega(|\mathrm{supp}(\boldsymbol{s})|)$ coordinates equal to -1. Thus, at line (7), $|\mathrm{supp}(\boldsymbol{s})|$ is reduced by a factor of $\kappa < 1$. It follows that the algorithm terminates after $O(\log(\frac{\varepsilon^2 m}{n})) = O(\log m)$ loop iterations. Further, at each iteration, we add $\sum_{i=1}^m x_i^* s_i\boldsymbol{A}_i$ to the matrix $\sum_{i=1}^m s_i\boldsymbol{A}_i$, which is originally $\boldsymbol{I}_n$. So by triangle inequality, at the end of the algorithm we have an additive error of at most

$$C \cdot \sum_{t\geq 0}\sqrt{\frac{n}{\kappa^t m}} = O\Big(\sqrt{\frac{n}{m}}\Big) = O(\varepsilon),$$

that is, $(1 - O(\varepsilon))\boldsymbol{I}_n \preceq \sum_{i=1}^m s_i\boldsymbol{A}_i \preceq (1 + O(\varepsilon))\boldsymbol{I}_n$. The error probability is dominated by $2^{-\Theta(m_0)}$, where $m_0 \geq n$ is the support in the last iteration. $\square$

6.5 Missing Proofs for Preliminaries

Proof of Cor. 109. Since $\mathbb{E}[\|\boldsymbol{y}\|_2] \le \mathbb{E}[\|\boldsymbol{y}\|_2^2]^{1/2} = \sqrt{m}$, we apply Theorem 4 to get, for $m \ge 7$,

$$\Pr_{\boldsymbol{y}\sim N(\boldsymbol{0},\boldsymbol{I}_m)}\left[\|\boldsymbol{y}\|_2 > m\right] \le e^{-(m-\sqrt{m})^2/2} \le 2^{-m}.$$

Since the function $y \mapsto -\frac{y^2}{2}+\log(y^2)$ is concave, we can upper bound it with any tangent line; in particular,

$$-\frac{y^2}{2} + \log(y^2) \le \left(\frac{2}{m} - m\right)\cdot y + \frac{m^2}{2} + \log(m^2) - 2,$$

so that using the standard estimate $P_{y\sim N(0,1)}[y > m] \ge \frac{m}{m^2+1}\cdot\frac{1}{\sqrt{2\pi}}e^{-m^2/2}$, we have

$$\mathop{\mathbb{E}}_{y\sim N(0,1)}[y^2 \mid y > m] \le \frac{\int_m^\infty \exp\left(\left(\frac{2}{m} - m\right)\cdot y + \frac{m^2}{2} + \log(m^2) - 2\right)dy}{\sqrt{2\pi}P[y > m]} \le \frac{m^2+1}{m}\cdot\frac{m^3}{m^2-2}$$

and therefore, for $m \ge 7$,

$$\mathbb{E}[\|\boldsymbol{y}\|_2^2 \mid \|\boldsymbol{y}\|_2 > m] \le m \cdot \mathop{\mathbb{E}}_{y\sim N(0,1)}[y^2 \mid y > m] < 2m^3.$$

Now, since

$$m = \mathbb{E}\left[\|\boldsymbol{y}\|_2^2\right] = \underbrace{\Pr\left[\|\boldsymbol{y}\|_2 > m\right]}_{\le \exp(-(m-\sqrt{m})^2/2)}\cdot\underbrace{\mathbb{E}\left[\|\boldsymbol{y}\|_2^2 \mid \|\boldsymbol{y}\|_2 > m\right]}_{\le 2m^3} + \underbrace{\Pr\left[\|\boldsymbol{y}\|_2 \le m\right]}_{\le 1}\cdot\mathbb{E}\left[\|\boldsymbol{y}\|_2^2 \mid \|\boldsymbol{y}\|_2 \le m\right],$$

it follows that $\mathbb{E}\left[\|\boldsymbol{y}\|_2^2 \mid \|\boldsymbol{y}\| \le m\right] \ge (1 - 2^{-m})\cdot m$ for $m \ge 7$. $\square$

Chapter 7

MATRIX BALANCING II: PARTIAL COLORING BOUNDS VIA MIRROR DESCENT

This chapter is based on a joint paper with Daniel Dadush and Haotian Jiang [52].

7.1 Introduction

Discrepancy minimization has been a well-studied area of research both in mathematics and computer science [45, 111]. We start with a classical setting: given vectors $a_1, \dots, a_n \in \mathbb{R}^d$ each satisfying $\|a_i\|_\infty \leq 1$, the goal is to find a coloring $x \in \{\pm 1\}^n$ that minimizes the discrepancy, defined as $\|\sum_{i=1}^n x_i a_i\|_\infty$. A seminal result of Spencer [157] improves upon the $O(\sqrt{n \log d})$ bound of a random coloring via Chernoff and union bound:

Theorem 117 (Spencer [157])**.** *Let $d \geq n$. Given vectors $a_1, \dots, a_n \in \mathbb{R}^d$ with $\|a_i\|_\infty \leq 1$, there exists $x \in \{\pm 1\}^n$ such that $\|\sum_{i=1}^n x_i a_i\|_\infty \lesssim \sqrt{n \log(2d/n)}$.*

In particular, when $d = n$, Theorem 117 states that the discrepancy is at most $O(\sqrt{n})$, as opposed to the $O(\sqrt{n \log n})$ bound for a random coloring. Spencer's theorem is known to be tight up to constants for all $d \geq n$ [45, 111].

The Partial Coloring Method. All known proofs of Spencer's theorem are essentially based on the *partial coloring* method, one of the most important and widely applied techniques in discrepancy theory. The method states that to obtain the type of discrepancy bound in Theorem 117, it suffices to prove the same bound for a partial coloring $x \in [-1, 1]^n$ with at least $\Omega(n)$ coordinates in $\{\pm 1\}$. This process is then iterated over the set of coordinates $\{i : |x_i| < 1\}$ to obtain a full coloring. For Spencer-type problems, the discrepancy of the full coloring is at most a constant factor off from the discrepancy of the partial coloring (see Corollary 8).

The partial coloring method was developed in the early 80s by Beck and refined by Spencer using the entropy method [32, 157]. A convex geometry view of partial coloring was developed independently by Gluskin [68]. While these original arguments used the pigeonhole principle and were non-algorithmic, a breakthrough result of Bansal [16], followed by a rich line of work [103, 146, 100, 61, 140], gave various algorithmic versions. These recent developments also led to new results in approximation algorithms and differential privacy [145, 126, 17, 25].

Matrix Spencer Setting. A natural generalization of Spencer's setting to matrices is the following. Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$, each satisfying $\|A_i\|_{\mathrm{op}} \leq 1$, the goal is to find a coloring $x \in \{\pm 1\}^n$ that minimizes $\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}}$. In particular, Spencer's setting corresponds to the case where all matrices A_i are diagonal.

In the matrix Spencer setting, the non-commutative Khintchine inequality of Lust-Piquard and Pisier [104, 134] shows that a random coloring $x \in \{\pm 1\}^n$ has expected discrepancy $\mathbb{E}[\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}}] \lesssim \sqrt{n \log r}$, where each matrix A_i has rank at most $r \leq d$. It is conjectured that the discrepancy bound in Theorem 117 can be generalized as follows:

Conjecture 14 (Matrix Spencer Conjecture [113, 171])**.** *Let $d \geq \sqrt{n}$. Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{\mathrm{op}} \leq 1$, there exists $x \in \{\pm 1\}^n$ such that*

$$\Big\|\sum_{i=1}^n x_i A_i\Big\|_{\mathrm{op}} \lesssim \sqrt{n \cdot \max(1, \log(d/n))}.$$

In particular, when $\sqrt{n} \leq d \leq n$, the conjectured discrepancy bound is $O(\sqrt{n})$. Despite significant effort, Conjecture 14 has remained largely open, with partial progress for block diagonal matrices [100]. A subsequent work showed that Conjecture 14 holds for matrices of rank at most $n/\log^3(n)$, and that in general a bound of $O(n^{1/4} d^{1/4} \log d)$ holds [21].

We also note that the condition $n \leq d^2$ is justified by Theorem 3, since for $n > d^2$ the bound is only worse by a factor of two.

Matrix Discrepancy for Schatten Norms. More generally, let[1] $2 \leq p \leq q \leq \infty$, we consider

[1]We make the assumption that $p \leq q$ to avoid a polynomial dependence on d in the discrepancy bound.

the following matrix discrepancy setting for Schatten norms. Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$, each satisfying $\|A_i\|_{S_p} \leq 1$, where $\|\cdot\|_{S_p}$ denotes the Schatten-p norm. The goal is to find a coloring $x \in \{\pm 1\}^n$ to minimize $\|\sum_{i=1}^n x_i A_i\|_{S_q}$, the $S_p \to S_q$ discrepancy. In particular, the matrix Spencer setting corresponds to the case where $p = q = \infty$.

The diagonal case of $S_p \to S_q$ discrepancy, i.e. $\ell_p \to \ell_q$ discrepancy for vectors, is well studied (see [53, 140] and the references therein). In fact, the well-known Komlós conjecture asserts that the $\ell_2 \to \ell_\infty$ discrepancy can be upper bounded by a universal constant. For general $\ell_p \to \ell_q$ discrepancy, Reis and Rothvoss [140] proves an optimal partial coloring bound of $O(\sqrt{\min(p, \log(d/n))} \cdot n^{1/2-1/p+1/q})$, assuming $d \geq n$ and $2 \leq p \leq q \leq \infty$. It is a natural question whether these bounds generalize to $S_p \to S_q$ discrepancy.

The Challenge in Using Partial Coloring Method for Matrix Discrepancy. Central to the partial coloring method is to show that the discrepancy body $D := \{x \in \mathbb{R}^n : \|\sum_{i=1}^n x_i A_i\| \leq t\}$, i.e. the set of fractional colorings with discrepancy at most t under norm $\|\cdot\|$, is "large" in some sense. A natural notion of largeness, due to Gluskin [68], is that the body D has Gaussian measure at least $2^{-O(n)}$. This measure of largeness has been adopted (sometimes implicitly) in essentially all work on partial coloring [16, 103, 146, 100, 61, 140].

For the setting in Theorem 117, the discrepancy body D is a polytope defined by the intersection of strips of the form $|\langle r_i, x\rangle| \leq t$, where $r_i \in \mathbb{R}^n$ are the rows of the $d \times n$ matix whose columns are $a_1, \ldots, a_n$. Therefore, Šidák's lemma [155] can be readily used to give a Gaussian measure lower bound of the form $\gamma_n(D) \geq \prod_{i=1}^d \gamma_n(\{x \in \mathbb{R}^n : |\langle r_i, x\rangle| \leq t\})$.

In the setting of matrix discrepancy, however, the discrepancy body D has an infinite number of facets. This prevents the use of Gaussian correlation inequalities to lower bound $\gamma_n(D)$. To get around this barrier and use the partial coloring method for matrix discrepancy, one needs a different approach for proving Gaussian measure lower bounds.

If $q < p$, then even a single matrix (i.e. $n = 1$) can have discrepancy $d^{1/q-1/p}$.

7.1.1 Our Results

We lower bound the Gaussian measure of the discrepancy body D via covering numbers for its polar D° with respect to the ℓ_∞-ball (see Section 7.3.1). We then prove the desired covering number estimates using mirror descent, the powerful convex optimization primitive of Nemirovski and Yudin [121] (see Sections 7.3.2 to 7.3.4). Our method yields the following applications.

Matrix Spencer for Low-Rank Matrices. Our first result is the following improvement over the $O(\sqrt{n \log r})$ bound for random coloring in the matrix Spencer setting.

Theorem 118 (Matrix Spencer for Low-Rank Matrices)**.** *Let $d \geq \sqrt{n}$. Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{\mathrm{op}} \leq 1$ and $\mathrm{rank}(A_i) \leq r$ for all $i \in [n]$, one can efficiently find a coloring $x \in \{\pm 1\}^n$ such that*

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_{\mathrm{op}} \lesssim \sqrt{n \cdot \max(1, \log(r \cdot \min(1, d/n)))}.$$

When the input matrices have rank $r \lesssim n/d$, the discrepancy bound in Theorem 118 is $O(\sqrt{n})$ and this proves Conjecture 14 for low rank matrices in the regime where $d \leq n$. Recall that this is subsumed by the main result of [21].

Matrix Spencer for Block Diagonal Matrices. Our second application is the following improved matrix Spencer bound for block diagonal matrices. Unlike the theorem above, this is not implied by [21].

Theorem 119 (Matrix Spencer for Block Diagonal Matrices)**.** *Let $d \geq \sqrt{n}$ and $h \leq d$. Given block diagonal symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{\mathrm{op}} \leq 1$ and block size $h \times h$, one can efficiently find a coloring $x \in \{\pm 1\}^n$ with*

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_{\mathrm{op}} \lesssim \sqrt{n \cdot \max(1, \log(hd/n))}.$$

In particular, Theorem 119 proves Conjecture 14 whenever $h \lesssim n/d$. This bound was previously proved in [100] under the assumption $h \leq \sqrt{n}$, which we remove here.

We also obtain the following reduction of Conjecture 14 to the construction of a better quantum relative entropy net for the spectraplex $\mathcal{S}_d := \{X \in \mathbb{R}^{d\times d} : X \succeq 0, \mathrm{Tr}(X) = 1\}$.

Corollary 120 (Better Entropy Net Implies Matrix Spencer). *Let $d \geq \sqrt{n}$. If we can find $T \subseteq \mathcal{S}_d$ with $|T| \leq 2^{O(n)}$ such that for each $X \in \mathcal{S}_d$ there exists $Y \in T$ with $S(X\|Y) \lesssim \max(1, \log(d/n))$, where $S(X\|Y)$ is the quantum relative entropy between X and Y, then the matrix Spencer conjecture is true.*

In particular, in the proof of Theorem 119, we construct a $O(\max(1, \log(hd/n)))$-relative entropy net for the set of block diagonal matrices on $\mathcal{S}_d$ with block size $h \times h$ (see Section 7.3.4). Our construction of such relative entropy nets might be of independent interest.

Matrix Discrepancy for Schatten Norms. Theorem 118 is a special case of the following general matrix discrepancy bound for Schatten norms.

Theorem 121 (Matrix Discrepancy for Schatten Norms). *Let $d \geq \sqrt{n}$ and $2 \leq p \leq q \leq \infty$. Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$ and $\mathrm{rank}(A_i) \leq r$ for all $i \in [n]$, one can efficiently find $x \in [-1, 1]^n$ so that $|\{i : |x_i| = 1\}| \geq n/2$ and*

$$\Big\| \sum_{i=1}^{n} x_i A_i \Big\|_{S_q} \lesssim \sqrt{n \cdot \min(p, \max(1, \log(rk)))} \cdot k^{1/p-1/q},$$

where we denote $k := \min(1, d/n)$. Moreover, we can find a full coloring $x \in \{\pm 1\}^n$ at the expense of a factor of $(1/2 + 1/q - 1/p)^{-1}$.

Our partial coloring result in Theorem 121 is tight when either $d = \Theta(\sqrt{n})$ (for which we give an alternative proof using Banaszczyk's result [13] in Section 7.7), or when $r = 1$ and $d \geq n$. We provide matching lower bounds for both cases in Sections 7.6.1 and 7.6.2. In particular, our lower bound examples imply a tight $\Omega(\sqrt{n})$ lower bound for rank-1 matrix Spencer when $d = n$.

Corollary 122 (Rank-1 Matrix Spencer Lower Bound). *There exist rank-1 symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{n\times n}$ with $\|A_i\|_{\mathrm{op}} \leq 1$ such that any $x \in \{\pm 1\}^n$ has $\|\sum_{i=1}^{n} x_i A_i\|_{\mathrm{op}} \gtrsim \sqrt{n}$.*

Another immediate consequence of our lower bounds is an optimal $\Omega(\sqrt{\min(d,n)})$ lower bound for $S_2 \to S_\infty$ discrepancy. This is in stark contrast to the well-known Komlós conjecture for vectors, which asserts that the $\ell_2 \to \ell_\infty$ discrepancy is $O(1)$. Corollary 123 states that such a conjecture is far from being true for matrices.

Corollary 123 (Lower Bound for Matrix Komlós). *For any d and n, there exist symmetric matrices $A_1, \dots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_F \le 1$ such that any $x \in \{\pm 1\}^n$ has $\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}} \gtrsim \sqrt{\min(d,n)}$.*

Finally, we propose the following generalization of Conjecture 14:

Conjecture 15 ($S_p \to S_q$ Matrix Discrepancy). *Let $d \ge \sqrt{n}$ and $2 \le p \le q \le \infty$. Given matrices $A_1, \dots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \le 1$, there exists $x \in \{\pm 1\}^n$ such that*

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_{S_q} \lesssim \sqrt{n \cdot \min(p, \max(1, \log(d/n)))} \cdot \min(1, d/n)^{1/p-1/q}.$$

When $d = n$, the right hand side is $O(\sqrt{n})$, and for diagonal matrices the conjecture is known to be true for any $2 \le p \le q$. When $p = q$, the conjecture is also known to be true for diagonal matrices for all d and n [140].

7.1.2 Overview of Our Approach

We give a brief overview of our partial coloring framework in this subsection, and leave a more detailed discussion to Section 7.3.

Partial Coloring via Covering Numbers. Let $K := \{x \in \mathbb{R}^n : \|\sum_{i=1}^n x_i A_i\| \le 1\}$ be the unit discrepancy body[2] and t be the target discrepancy bound. A recent refinement by Reis and Rothvoss [140] of Gluskin's convex geometry approach [68] shows that whenever $\gamma_n(tK) \ge 2^{-O(n)}$ for any constant in the exponent, one can efficiently find a partial coloring $x \in O(tK) \cap [-1,1]^n$ with at least $n/2$ coordinates in $\{-1, 1\}$ (see Theorem 7). For settings

[2] To avoid confusion when talking about discrepancy bodies, K denotes the unit discrepancy body, and D denotes a scaling of K by the target discrepancy bound.

where the target discrepancy bound is $n^{\Omega(1)}$, we may iterate the partial coloring to find a full coloring with the same discrepancy bound up to constants (see Corollary 8).

Our new approach for proving a Gaussian measure lower bound $\gamma_n(tK) \geq 2^{-O(n)}$ is via the covering numbers (Definition 36) of K or K° with respect to the Euclidean ball B_2^n or the ℓ_∞ ball B_∞^n. In particular, since $\gamma_n(\sqrt{n}B_2^n)$ has constant Gaussian measure, as long as $\mathcal{N}(\sqrt{n}B_2^n, tK) \leq 2^{O(n)}$, we get $\gamma_n(tK) \geq 2^{-O(n)}$. Using the duality of covering numbers and connections with volume, we obtain several equivalent conditions for $\gamma_n(tK) \geq 2^{-O(n)}$ in terms of covering (Lemma 46). The condition that we will work with is $\mathcal{N}(K^\circ, \frac{t}{n}B_\infty^n) \leq 2^{O(n)}$, where $K^\circ = \{(\langle A_1, U\rangle, \ldots, \langle A_n, U\rangle) : \|U\|_* \leq 1\}$ is the polar discrepancy body.

Covering via Mirror Descent. We prove the covering number bound $\mathcal{N}(K^\circ, \frac{t}{n}B_\infty^n) \leq 2^{O(n)}$ using mirror descent, a powerful convex optimization primitive of Nemirovski and Yudin [121] (see Section 7.3.2 for an overview). In particular, denote the linear map $\mathcal{A}(U) := (\langle A_1, U\rangle, \ldots, \langle A_n, U\rangle)$. We shall assume that each $\|A_i\| \leq 1$. This is true for the matrix Spencer setting with $\|\cdot\|$ being the operator norm. In the more general setting of matrix discrepancy for Schatten norms, we have $\|A_i\|_{S_p} \leq 1$ while the norm for measuring discrepancy is $\|\cdot\|_{S_q}$. One can get around this issue by leveraging known covering number estimates between Schatten classes (Theorem 125).

For any matrix $\|U\|_* \leq 1$, consider minimizing the function $f_U(X) := \|\mathcal{A}(X - U)\|_\infty$ over the dual unit ball $B_* := \{U : \|U\|_* \leq 1\}$. The function has minimum value $f_U(U) = 0$ and since it has subgradients in $\{\pm A_1, \ldots, \pm A_n\}$ with $\|A_i\| \leq 1$, the function $f_U(X)$ is 1-Lipschitz with respect to the dual norm $\|\cdot\|_*$. So as long as there exists a 1-strongly convex mirror map Φ on B_*, we can minimize $f_U(X)$ by starting from some matrix $U_0 = U_0(U) \in B_*$ and running mirror descent for n steps. Denoting by U_s the matrix in the s-th step, standard guarantees for mirror descent (Theorem 127) yield

$$\min_{s\in[n]} f_U(U_s) = \min_{s\in[n]} f_U(U_s) - f_U(U) \leq \sqrt{\frac{2D_\Phi(U, U_0)}{n}}, \tag{7.1}$$

where $D_\Phi(U, U_0) = \Phi(U) - \Phi(U_0) - \langle \nabla\Phi(U_0), U - U_0\rangle$ is the Bregman divergence. We let T be the set of all matrices encountered when running mirror descent for all possible $U \in B_*$, i.e. $T := \{U_s : s \in [n], U \in B_*\}$, and $T_0 := \{U_0 : U \in B_*\}$ be the set of all starting matrices. The net $\mathcal{A}(T)$ will be our covering for K°.

To see that this indeed gives a good covering, we denote $D_\Phi^{\max} := \sup_{U \in B_*} D_\Phi(U \| U_0)$. By the definition of the function f_U, we have from (7.1) that

$$\min_{s\in[n]} \|\mathcal{A}(U) - \mathcal{A}(U_s)\|_\infty \le \sqrt{\frac{2D_\Phi(U, U_0)}{n}} \le \sqrt{\frac{2D_\Phi^{\max}}{n}},$$

and so the dual body admits the covering $K^\circ \subseteq \mathcal{A}(T) + \sqrt{2D_\Phi^{\max}/n} \cdot B_\infty^n$. Thus as long as our target discrepancy bound $t \le \sqrt{2nD_\Phi^{\max}}$, we have $\mathcal{N}(K^\circ, \frac{t}{n} B_\infty^n) \le |T|$, which we need to show to be at most $2^{O(n)}$.

The key observation we make here is that for our choices of the mirror maps in Sections 7.4 and 7.5, U_s only depends[3] on the sum of the subgradients, but not on their order. Since there are only $2n$ choices of subgradients $\{\pm A_i\}_{i\in[n]}$ and we run mirror descent for n steps, a counting argument reveals that there are at most $2^{O(n)}$ possible sums of gradients (Lemma 128). So long as the starting matrices satisfy $|T_0| \le 2^{O(n)}$, we have $|T| \le |T_0| \cdot 2^{O(n)} \le 2^{O(n)}$.

A View of Mirror Descent as Refining the Net. In the diagonal case, i.e. Spencer's setting, we can directly build the net T by repeatedly sampling the ith diagonal coordinate $e_ie_i^\top$ proportional to its weight in the target matrix. Since the set of diagonal matrices on the Schatten-1 ball has only $2d$ vertices $\{\pm e_ie_i^\top\}_{i\in[d]}$, the approximate Carathéodory theorem (see [167], Theorem 0.0.2) implies that the image of the net $\mathcal{A}(T)$ already gives a good covering for K°, and mirror descent is not necessary in this case.

However, this argument fails beyond diagonal matrices, as the number of vertices becomes infinite. In these more general cases, we use mirror descent to boost a coarse net T_0 to a finer net T which has a better covering guarantee in the image space, at the expense of increasing the size of the net by a factor of $2^{O(n)}$.

[3]In general, mirror descent projects back onto the feasible set according to the Bregman divergence in each iteration, and therefore might not satisfy this property.

Relative Entropy Nets for the Spectraplex. For our application in Section 7.5 to low-rank matrices, it suffices to take $T_0 = \{0\}$. For the application in Section 7.4 to block diagonal matrix Spencer, we run mirror descent on the spectraplex $\mathcal{S}_d := \{X \in \mathbb{R}^{d \times d} : X \succeq 0, \mathrm{Tr}(X) = 1\}$ and carefully construct a set $|T_0| \leq 2^{O(n)}$ with small $D_\Phi^{\max}$. Since $D_\Phi(X\|Y)$ is the quantum relative entropy between X and Y in the spectraplex setup, we refer to such T_0 as a (quantum) relative entropy net (Definition 129).

We use an operator norm net for the Schatten-1 ball from [73] to construct a relative entropy net with error $O(\log(d^2/n))$ for the spectraplex $\mathcal{S}_d$ (Lemma 130). When restricted to block diagonal matrices with block size $h \times h$, we use a hybrid of this argument and the earlier approximate Caratheodory argument to find a refined relative entropy net with error $O(\log(hd/n))$ (Theorem 131). Taking T_0 to be this net in our mirror descent framework gives Theorem 119. This also allows us to reduce the matrix Spencer conjecture to the existence of a better relative entropy net with error $O(\log(d/n))$ for the spectraplex (Corollary 120).

7.1.3 *Further Related Work*

Banaszczyk's Approach. While the partial coloring method has been extensively applied in discrepancy and obtains the optimal bound for many problems, for several applications where the target discrepancy bound is $n^{o(1)}$ (e.g. the Komlós problem or Tusnady's problem), partial coloring is potentially sub-optimal by a logarithmic factor. In breakthrough work, Banaszczyk [13] obtained an improvement over the partial coloring method for these applications using deep techniques from convex geometry. While Banaszczyk's original proof is non-constructive, a fascinating recent line of work has obtained algorithmic versions of Banaszczyk's result [51, 18, 20, 100, 19].

Matrix Spencer Conjecture and Non-commutative Random Matrix Theory. The typical value of $\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}}$ for a random coloring has attracted significant attention in random matrix theory. For commutative matrices, the bound $\mathbb{E}[\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}}] \lesssim \sqrt{n \log d}$ by matrix Khintchine [104, 134] or matrix Chernoff bound [1] is in general tight. It is also known to be tight for Toeplitz matrices [112]. For matrices with certain non-commutative structures (e.g. random Gaussian matrices), improved bounds of $O(\sqrt{n})$ are known (see [167]). In the context of Conjecture 14, these results imply that a random coloring already achieves the conjectured bound when the input matrices have certain non-commutative structures. On the other hand, by Theorem 117, Conjecture 14 is known when all the matrices commute.

Concurrent and Independent Work. In concurrent and independent work, Hopkins, Raghavendra and Shetty [76] proved a bound of $\sqrt{n \log\left(\mathrm{Tr}(\sum_{i=1}^n A_i^2)/n^{1.5}\right)}$ for matrix Spencer using quantum communication complexity. Their bound coincides with ours for full rank matrices, and is slightly stronger for low-rank matrices. However, our approach is completely different and can also be used to show matrix discrepancy bounds for block diagonal matrices and general Schatten norms. We believe both approaches are interesting and may lead to further progress in resolving the matrix Spencer conjecture.

7.2 *Preliminaries*

Norms and Convex Bodies. A convex body is a compact convex set with non-empty interior. We say a convex set K is symmetric if $x \in K$ implies $-x \in K$. We use $\|\cdot\|_p$ to denote the ℓ_p-norm and $\|\cdot\|_{S_p}$ to denote the Schatten-p norm. In particular, the operator norm $\|\cdot\|_{\mathrm{op}} = \|\cdot\|_{S_\infty}$ and the Frobenius norm $\|\cdot\|_F = \|\cdot\|_{S_2}$. We use B_p^n to denote the unit ℓ_p-ball in $\mathbb{R}^n$ and $B_{S_p}^n := \{A \in \mathbb{R}^{n\times n} : \|A\|_{S_p} \le 1\}$ to denote the unit Schatten-p ball in $\mathbb{R}^{n\times n}$, with $B_{\mathrm{op}}^n := B_{S_\infty}^n$. Let $\mathbb{R}_+^n$ denote the set of non-negative vectors in $\mathbb{R}^n$ and denote the simplex $\Delta_n := \{x \in \mathbb{R}_+^n : \|x\|_1 = 1\}$. Let $\mathbb{S}_+^n$ (resp. $\mathbb{S}_{++}^n$) denote the set of positive semidefinite (resp. positive definite) $n \times n$ matrices, and define the spectraplex $\mathcal{S}_n := \{X \in \mathbb{S}_+^n : \mathrm{Tr}(X) = 1\}$. For a norm $\|\cdot\|$ in $\mathbb{R}^n$, we define the dual norm as $\|x\|_* := \sup\{\langle y, x\rangle : y \in \mathbb{R}^n, \|y\| \le 1\}$. Dual norms are similarly defined for matrix norms.

Convex Functions. A convex function $f : \mathcal{X} \to \mathbb{R}$ is said to be L-Lipschitz with respect to a norm $\|\cdot\|$ if $\|g\|_* \leq L$ for all subgradients $g \in \partial f(x)$. We say that f is α-strongly convex with respect to a norm $\|\cdot\|$ if $f(y) \geq f(x) + g^\top(y-x) + \frac{\alpha}{2}\|x-y\|^2$, for all $x, y \in \mathcal{X}$ and all subgradients $g \in \partial f(x)$.

Polar. Given a convex set $K \subseteq \mathbb{R}^n$ with $0 \in K$, we define the polar of K to be $K^\circ := \{y \in \mathbb{R}^n : \sup_{x\in K}\langle x, y\rangle \leq 1\}$. It is immediate from the definition that for any constant $t > 0$, $(tK)^\circ = \frac{1}{t}K^\circ$. When K is closed, the polarity theorem states that $(K^\circ)^\circ = K$.

Lemma 124 (Polar of Discrepancy Set)**.** *Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ and a norm $\|\cdot\|$ in $\mathbb{R}^{d\times d}$, we define the unit discrepancy set as $K := \{x \in \mathbb{R}^n : \|\sum_{i=1}^n x_i A_i\| \leq 1\}$. Then $K' := \{(\langle A_1, U\rangle, \ldots, \langle A_n, U\rangle) : \|U\|_* \leq 1\}$ is the polar body $K' = K^\circ$.*

Proof. By the definition of polar body, we may write

$$\begin{aligned}(K')^\circ &= \Big\{x \in \mathbb{R}^n : \sum_{i=1}^n x_i\langle A_i, U\rangle \leq 1,\ \forall\, U \text{ s.t. } \|U\|_* \leq 1\Big\} \\ &= \Big\{x \in \mathbb{R}^n : \Big\langle \sum_{i=1}^n x_i A_i, U\Big\rangle \leq 1,\ \forall\, U \text{ s.t. } \|U\|_* \leq 1\Big\} \\ &= K,\end{aligned}$$

by the definition of dual norm. It then follows from the polarity theorem that $K' = K^\circ$. $\square$

We will also need the following upper bound on the covering numbers of Schatten balls[4].

Theorem 125 ([73], Theorem 1.1)**.** *Let $d, n \in \mathbb{N}$ and $1 \leq p \leq q \leq \infty$. Then we have*

$$\mathcal{N}\Big(B^d_{S_p}, \min\Big(1, \frac{d}{n}\Big)^{1/p-1/q} B^d_{S_q}\Big) \leq 2^{O(n)}.$$

[4]We note that [73] claims the bound only up to a constant depending on p and q, but their argument readily gives a universal constant in the regime of $p, q \geq 1$.

7.3 *Our Framework for Partial Coloring*

7.3.1 *Partial Coloring via Covering Numbers*

Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$, a norm $\|\cdot\|$ on $\mathbb{R}^{d\times d}$ for measuring the discrepancy, and a target discrepancy bound t, let $D := \{x \in \mathbb{R}^n : \|\sum_{i=1}^n x_i A_i\| \leq t\}$ be the associated discrepancy body. For our mirror descent framework, we use the following:

Corollary 126. *Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$, let $K_{q+}^\circ := \{(\langle A_1, U\rangle, \ldots, \langle A_n, U\rangle) : U \in B_{S_q}^d, U \succeq 0\}$. If we have $\mathcal{N}(K_{q+}^\circ, \frac{t}{n}B_\infty^n) \leq 2^{O(n)}$, then we can efficiently find a partial coloring $x \in [-1,1]^n$ with $|\{i : |x_i| = 1\}| \geq n/2$ and $\|\sum_{i=1}^n x_i A_i\|_{S_q} \lesssim t$.*

Proof. Since $tD^\circ \subseteq K_{q+}^\circ - K_{q+}^\circ$, by Lemma 38 we have $\mathcal{N}(D^\circ, \frac{1}{n}B_\infty^n) = 2^{O(n)}$. The equivalence $(1) \Leftrightarrow (5)$ in Lemma 46 implies $\gamma_n(D) \geq 2^{-O(n)}$, and Theorem 7 gives the corollary. $\square$

7.3.2 *Mirror Descent: An Overview*

The mirror descent method was introduced by Nemirovski and Yudin [121]. Here, we follow the presentation in [39]. Let $\mathcal{D}$ be an open subset of $\mathbb{R}^d$ and $\mathcal{X}$ a subset of its closure. We fix a convex function $f : \mathcal{X} \to \mathbb{R}$ assumed to be L-Lipschitz with respect to a norm $\|\cdot\|$, and a differentiable function $\Phi : \mathcal{D} \to \mathbb{R}$ that is ρ-strongly convex with respect to $\|\cdot\|$ and has a surjective gradient $\nabla\Phi : \mathcal{D} \to \mathbb{R}^d$. The mirror descent algorithm, given a starting point $x_0 \in \mathcal{X} \cap \mathcal{D}$, consists of the iterations

$$\nabla\Phi(y_{t+1}) := \nabla\Phi(x_t) - \eta g_t,$$

$$x_{t+1} := \operatorname{argmin}_{x\in\mathcal{X}\cap\mathcal{D}} D_\Phi(x, y_{t+1}),$$

where $g_t \in \partial f(x_t)$ and $D_\Phi(x,y) := \Phi(x) - \Phi(y) - \nabla\Phi(y)^\top(x-y)$ is the Bregman divergence. Note that $y_t \in \mathcal{D}$ and $x_t \in \mathcal{X}\cap\mathcal{D}$ for all $t \geq 0$. We use the following convergence guarantee:

Theorem 127 ([39], Theorem 4.2)**.** *Let f be L-Lipschitz and Φ be ρ-strongly convex with respect to $\|\cdot\|$, and $D_\Phi^{\max} \geq D_\Phi(x^*, x_0)$ be any upper bound. Then the mirror descent algorithm with*

$\eta := \frac{1}{L}\sqrt{\frac{2\rho D_\Phi^{\max}}{T}}$ *satisfies*

$$\min_{s\in[T]} f(x_s) - f(x^*) \leq L\sqrt{\frac{2D_\Phi^{\max}}{\rho T}}.$$

The Spectraplex Setup. Here we take $\mathcal{X} := \mathcal{S}_d = \{X \in \mathbb{S}_+^d : \mathrm{Tr}(X) = 1\}$. The mirror map is $\Phi(X) = \mathrm{Tr}(X \log X)$, defined on $\mathcal{D} = \mathbb{S}_{++}^d$, which is $\frac{1}{2}$-strongly convex with respect to the Schatten-1 norm by the quantum Pinsker inequality [42]. Then the convergence bound in Theorem 127 becomes $2L\sqrt{\frac{S(X^*\|X_0)}{T}}$, where $S(X\|Y) := \mathrm{Tr}(X(\log X - \log Y))$ is the quantum relative entropy between matrices $X, Y \in \mathcal{S}_d$. The projection step corresponds to a trace normalization, so given a starting point $X_0 \in \mathcal{S}_d \cap \mathbb{S}_{++}^d$, we may write in closed form

$$X_{t+1} = \frac{\exp(\log X_0 - \eta \sum_{i=0}^t g_i)}{\mathrm{Tr}(\exp(\log X_0 - \eta \sum_{i=0}^t g_i))}, \tag{7.2}$$

for subgradients $g_i \in \partial f(X_i)$.

The Schatten Norm Setup. Here we take $\mathcal{X} = \mathcal{D} = \mathbb{R}^{d\times d}$, so that $X_t = Y_t$ for all t. The mirror map is $\Phi(X) := \frac{1}{2(p-1)}\|X\|_p^2$, which is known to be 1-strongly convex for all $p \in (1, 2]$ [9]. Thus given a starting point $X_0 \in \mathbb{R}^{d\times d}$, we may write in closed form

$$X_{t+1} = \nabla\Phi^{-1}\Big(\nabla\Phi(X_0) - \eta\sum_{i=0}^t g_i\Big), \tag{7.3}$$

for subgradients $g_i \in \partial f(X_i)$.

7.3.3 *Covering via Mirror Descent*

Given symmetric matrices $A_1, \ldots, A_n$ with $\|A_i\| \leq 1$ for all $i \in [n]$, where the dual norm $\|\cdot\|_*$ is either the Schatten-1 norm or the Schatten-p norm for some $p \in (1, 2]$, we apply mirror descent on functions of the form $f_U(X) := \max_{i\in[n]} |\langle A_i, X - U\rangle|$ to cover the polar discrepancy body

$$K^\circ := \{\mathcal{A}(U) : \|U\|_* \leq 1\}, \text{ where } \mathcal{A}(U) := (\langle A_1, U\rangle, \ldots, \langle A_n, U\rangle).$$

Note that $f_U(X) = \|\mathcal{A}(X) - \mathcal{A}(U)\|_\infty$ and that f is 1-Lipschitz with respect to $\|\cdot\|_*$. The key property of such functions is that we may always choose subgradients from the set of $2n$ matrices $\{\pm A_i : i \in [n]\}$, which allows us to upper bound the number of different matrices encountered during the mirror descent process.

Lemma 128. *Let $\|\cdot\|_*$ be either $\|\cdot\|_{S_1}$ as in the Spectraplex Setup, or $\|\cdot\|_{S_p}$ with $p \in (1,2]$ as in the Schatten Norm Setup, and $\mathcal{X}, \mathcal{D}$ be defined accordingly. Let $T_0 \subseteq \mathcal{X} \cap \mathcal{D}$ be a set with size $|T_0| \leq 2^{O(n)}$ and $K^\circ \supseteq K' = \mathcal{A}(T')$ the convex body to be covered, where $T' \subseteq \mathcal{X} \cap \mathcal{D}$. If for every $U \in T'$ there exists a starting point $U_0 := U_0(U) \in T_0$ with $D_\Phi(U, U_0) \leq D_\Phi^{\max}$, then we can bound*

$$\mathcal{N}\left(K', \sqrt{\frac{D_\Phi^{\max}}{n}} B_\infty^n\right) \leq 2^{O(n)}.$$

Proof. The key observation is that in either setup of mirror descent, the point X_t in (7.2) or (7.3) depends only on the starting point U_0 and on the sum of gradients $g_0, \ldots, g_{t-1}$, but not on their order. Moreover, we can always choose from the set of $2n$ gradients $\{\pm A_i : i \in [n]\}$ at each step. Thus applying mirror descent to the function f_U for all possible U with the same starting point U_0, the total number $N(U_0)$ of points visited in $T := n$ iterations satisfies

$$N(U_0) \leq \sum_{t=0}^{n} \binom{t + 2n - 1}{2n - 1} \leq (n+1) \cdot \binom{3n}{n} \leq 2^{O(n)}.$$

Since $|T_0| \leq 2^{O(n)}$, we obtain a set of $2^{O(n)}$ points $\mathcal{U}$ such that for every $Y = \mathcal{A}(U) \in K'$, there exists some $\tilde{U} \in \mathcal{U}$ so that $\|\mathcal{A}(\tilde{U}) - \mathcal{A}(U)\|_\infty = f_U(\tilde{U}) = f_U(\tilde{U}) - f_U(U) \lesssim \sqrt{D_\Phi^{\max}/n}$. $\square$

In the Schatten Norm Setup, we shall pick $K' = K^\circ$ and $T_0 = \{0\}$, i.e. U_0 is always 0. For the Spectraplex Setup, we carefully choose a set of starting points $|T_0| \leq 2^{O(n)}$ which has small $D_\Phi^{\max}$ with respect to $K' = \{\mathcal{A}(U) : U \in \mathcal{S}_d\}$. Since $D_\Phi(X\|Y)$ is the quantum relative entropy between X and Y in the Spectraplex Setup, we shall refer to the set of starting points T_0 as a (quantum) relative entropy net for $\mathcal{S}_m$.

Definition 129 (Quantum Relative Entropy Net)**.** *Given subsets $T, \mathcal{M} \subseteq \mathcal{S}_d$, T is a relative entropy net of $\mathcal{M}$ with error ε if for any $X \in \mathcal{M}$, we can find $Y \in T$ such that $S(X\|Y) \le \varepsilon$.*

7.3.4 Initialization for Spectraplex Setup: Relative Entropy Net

We start with the following lemma which constructs a relative entropy net on $\mathcal{S}_d$ from an operator norm net.

Lemma 130 (Relative Entropy Net from Operator Norm Net)**.** *Let $X, Y \in \mathcal{S}_d$ satisfies $\|X - Y\|_{\mathrm{op}} \le \varepsilon$ for some $\varepsilon \ge 1/d$. Then $S(X\|Y') \le \log(2d\varepsilon)$, where $Y' := \frac{1}{2}(Y + \frac{I_d}{d}) \in \mathcal{S}_d$.*

Proof. Recall that $\log(\cdot)$ is operator monotone and note that $X \preceq Y + \varepsilon I_d$. We then have

$$\begin{aligned} S(X\|Y') &= \mathrm{Tr}(X \cdot (\log X - \log Y')) \\ &\le \mathrm{Tr}(X \cdot (\log(Y + \varepsilon I_d) - \log Y')) \\ &\le \mathrm{Tr}(X) \cdot \|\log(Y + \varepsilon I_d) - \log Y'\|_{\mathrm{op}} \\ &\le \log\left(2 \cdot \left\|\frac{Y + \varepsilon I_d}{Y + \frac{I_d}{d}}\right\|_{\mathrm{op}}\right) \le \log(2d\varepsilon), \end{aligned}$$

where the first inequality follows from the operator monotonicity of $\log(\cdot)$, the second follows from matrix Hölder, and the last follows because $\varepsilon \ge 1/d$ and $\|Y\|_{\mathrm{op}} \le 1$. $\square$

Using the lemma above, we give the following construction for relative entropy nets on $\mathcal{S}_d$.

Theorem 131 (Entropy Net for Spectraplex)**.** *Given positive integers h, d and n such that d/h is an integer, let $\mathcal{S}_d^h \subseteq \mathcal{S}_d$ be the set of $d \times d$ block diagonal matrices on the spectraplex with block size $h \times h$. Then we can find a relative entropy net T for $\mathcal{S}_d^h$ with error at most $\max(1, \log(2hd/n))$ and size $|T| \le 2^{O(n)}$.*

Proof. By merging blocks as needed, we may assume $hd \ge n$. By Lemma 130, it suffices to find an operator norm net T' with size $|T'| \le 2^{O(n)}$ and distance $\varepsilon = \frac{\max\{h, \log(d/hn)\}}{n}$.

Let $\ell := d/h$ be the number of blocks, $X_1, \dots, X_\ell \in \mathbb{R}^{h \times h}$ denote the blocks of matrix $X \in \mathcal{S}_d^h$, and $N := 2/\varepsilon = 2n/\max\{h, \log(\ell/n)\}$ (we assume that N is an integer). Let $Z := \{z \in \mathbb{Z}_{\geq 0}^\ell : \sum_{i=1}^\ell z_i = N\}$, and for each $z \in Z$, we define

$$T_z := \{X \in \mathcal{S}_d^h : \mathrm{Tr}(X_i) = z_i/N, \forall i \in [\ell]\}.$$

It follows from a standard rounding argument that for any matrix $X \in \mathcal{S}_d^h$, one can find a matrix $Y \in \cup_{z \in Z} T_z$ with $\|X - Y\|_{\mathrm{op}} \leq 1/N = \varepsilon/2$.

We first show that $|Z| \leq 2^{O(n)}$. When $\ell \leq 2n$, we have

$$|Z| \leq \binom{N+\ell}{\ell} \leq \binom{N+2n}{2n} \leq \binom{\frac{2n}{h}+2n}{2n} \leq 2^{O(n)}.$$

When $\ell \geq 2n \geq N$, we can bound

$$|Z| \leq \binom{N+\ell}{N} \leq \binom{2\ell}{N} \leq \binom{2\ell}{\frac{2n}{\log(\ell/n)}} \leq \left(\frac{e\ell\log(\ell/n)}{n}\right)^{\frac{2n}{\log(\ell/n)}} \leq 2^{O(n)}.$$

It therefore suffices to construct an $\varepsilon/2$-operator norm net for each T_z.

Fix an arbitrary $z \in Z$. Note that the ith block of the matrices in T_z comes from $\frac{z_i}{N} \cdot \mathcal{S}_h$. Pick $n_i := z_i h$, we have from Theorem 125 that

$$\mathcal{N}\Big(\frac{z_i}{N}\mathcal{S}_h, \frac{z_i}{N} \cdot \frac{h}{n_i} B_{\mathrm{op}}^h\Big) = \mathcal{N}\Big(\mathcal{S}_h, \frac{h}{n_i} B_{\mathrm{op}}^h\Big) \leq 2^{O(n_i)}.$$

We denote this net as $\widetilde{T}_{z,i}$. It follows from the above that for any $X_i \in \frac{z_i}{N}\mathcal{S}_h$, there exists $Y_i \in \widetilde{T}_{z,i}$ with $\|X_i - Y_i\|_{\mathrm{op}} \leq \frac{z_i}{N} \cdot \frac{h}{n_i} = \varepsilon/2$. Define $\widetilde{T}_z := \{\mathrm{diag}(Y_1, \dots, Y_\ell) : Y_i \in \widetilde{T}_{z,i}\ \forall i \in [\ell]\}$. Then for any $X \in T_z$, there exists $Y \in \widetilde{T}_z$ such that $\|X - Y\|_{\mathrm{op}} \leq \varepsilon/2$, and thus $\widetilde{T}_z$ is indeed an $\varepsilon/2$-operator norm net for T_z. Furthermore, the size of $\widetilde{T}_z$ can be upper bounded as

$$|\widetilde{T}_z| \leq \prod_{i \in [\ell]} 2^{O(n_i)} = 2^{O(\sum_{i=1}^n z_i h)} = 2^{O(Nh)} \leq 2^{O(n)},$$

since $N \leq 2n/h$. This proves that $\widetilde{T} := \cup_{z \in Z} \widetilde{T}_z$ is an ε-operator norm net for $\mathcal{S}_d^h$ and has size at most $|\widetilde{T}| \leq 2^{O(n)}$, where we recall that $\varepsilon = \frac{\max\{h, \log(d/hn)\}}{n}$. Finally, invoking Lemma 130, $\widetilde{T}$ can be transformed into a relative entropy net T with size $|T| \leq 2^{O(n)}$ and error at most $\log(2d\varepsilon) \leq \log(2hd/n)$. This finishes the proof of the theorem. $\square$

7.4 Applications of the Spectraplex Setup

In this section, we prove our matrix Spencer bound for block diagonal matrices in Theorem 119, which we restate below.

Theorem 119 (Matrix Spencer for Block Diagonal Matrices)**.** *Let $d \geq \sqrt{n}$ and $h \leq d$. Given block diagonal symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d \times d}$ with $\|A_i\|_{\mathrm{op}} \leq 1$ and block size $h \times h$, one can efficiently find a coloring $x \in \{\pm 1\}^n$ with*

$$\Big\| \sum_{i=1}^{n} x_i A_i \Big\|_{\mathrm{op}} \lesssim \sqrt{n \cdot \max(1, \log(hd/n))}.$$

Proof of Theorem 119. By Theorem 131, we can find a relative entropy net T_0 of $\mathcal{S}_d^h$ with error $D_\Phi^{\max} := \max(1, \log(2hd/n))$ and size $|T_0| \leq 2^{O(n)}$. Then using Lemma 128 with the Spectraplex Setup for $K' := \mathcal{A}(\mathcal{S}_d^h)$ and T_0 being the relative entropy net, we obtain

$$\mathcal{N}\Big(K', \frac{t}{n} B_\infty^n\Big) \leq 2^{O(n)},$$

where $t = \sqrt{n \max(1, \log(2hd/n))}$. Let $\mathbb{S}_d^h$ be the set of $d \times d$ symmetric block diagonal matrices with block size $h \times h$. Define convex body $K'' := \mathcal{A}(B_{S_1}^d \cap \mathbb{S}_d^h \cap \mathbb{S}_+^d)$. We first prove that $\mathcal{N}(K'', \frac{t}{n} B_\infty^n) \leq 2^{O(n)}$. Since $\mathcal{N}(K', \frac{t}{n} B_\infty^n) \leq 2^{O(n)}$ by Theorem 131, we also have $\mathcal{N}(\frac{j}{n^2} K', \frac{t}{n} B_\infty^n) \leq 2^{O(n)}$ for each integer $j \in [n^2]$. We let H_j be the set of centers for the minimum covering of $\frac{j}{n^2} K'$ by translates of $\frac{t}{n} B_\infty^n$ and define $H = \cup_{j \in [n^2]} H_j$. Since $|H_j| \leq 2^{O(n)}$, it follows that $|H| \leq 2^{O(n)}$. For each $X \in B_{S_1}^d$ that satisfies $X \succeq 0$, we let $\frac{j}{n^2}$ be the multiple of $\frac{1}{n^2}$ that is closest to $\mathrm{Tr}(X)$, and set $X' := \frac{j}{n^2 \mathrm{Tr}(X)} \cdot X$. Then we have

$$\|\mathcal{A}(X') - \mathcal{A}(X)\|_\infty \leq \frac{1}{n^2} \cdot \|\mathcal{A}(X)\|_\infty \leq \frac{t}{n}.$$

As $\mathrm{Tr}(X') = \frac{j}{n^2}$, we can also find $Y \in H_j$ with $\|\mathcal{A}(X') - Y\|_\infty \leq \frac{t}{n}$. Therefore, $\|\mathcal{A}(X) - Y\|_\infty \leq \frac{2t}{n}$, and it follows that $K'' \subseteq H + \frac{2t}{n} B_\infty^n$. This implies $\mathcal{N}(K'', \frac{t}{n} B_\infty^n) \leq 2^{O(n)}$.

Next note that the dual discrepancy body $K^\circ := \mathcal{A}(B_{S_1}^d) = \mathcal{A}(B_{S_1}^d \cap \mathbb{S}_d^h)$ since each $A_i \in \mathbb{S}_d^h$. We have $K^\circ = K'' - K''$, so using Lemma 38 we get $\mathcal{N}(K^\circ, K'') \leq 2^{O(n)}$. Thus

$$\mathcal{N}\Big(K^\circ, \frac{t}{n} B_\infty^n\Big) \leq \mathcal{N}(K^\circ, K'') \cdot \mathcal{N}\Big(K'', \frac{t}{n} B_\infty^n\Big) \leq 2^{O(n)},$$

and $\gamma_n(tK) \geq 2^{-O(n)}$ by using Lemma 46. Corollary 8 then gives a full coloring $x \in \{\pm 1\}^n$ with discrepancy $\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}} \leq O(t)$. This finishes the proof of the theorem. $\square$

The analysis above also shows that if we can improve the bound in Theorem 131 to $O(\log(d/n))$ for any block size h, then the matrix Spencer conjecture is true.

Corollary 120 (Better Entropy Net Implies Matrix Spencer)**.** *Let $d \geq \sqrt{n}$. If we can find $T \subseteq \mathcal{S}_d$ with $|T| \leq 2^{O(n)}$ such that for each $X \in \mathcal{S}_d$ there exists $Y \in T$ with $S(X\|Y) \lesssim \max(1, \log(d/n))$, where $S(X\|Y)$ is the quantum relative entropy between X and Y, then the matrix Spencer conjecture is true.*

7.5 Matrix Discrepancy for Schatten Norms

In this section, we prove the following generalization of Theorem 118 for arbitrary Schatten norms by using a different regularizer for mirror descent.

Theorem 121 (Matrix Discrepancy for Schatten Norms)**.** *Let $d \geq \sqrt{n}$ and $2 \leq p \leq q \leq \infty$. Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$ and $\mathrm{rank}(A_i) \leq r$ for all $i \in [n]$, one can efficiently find $x \in [-1,1]^n$ so that $|\{i : |x_i| = 1\}| \geq n/2$ and*

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_{S_q} \lesssim \sqrt{n \cdot \min(p, \max(1, \log(rk)))} \cdot k^{1/p-1/q},$$

where we denote $k := \min(1, d/n)$. Moreover, we can find a full coloring $x \in \{\pm 1\}^n$ at the expense of a factor of $(1/2 + 1/q - 1/p)^{-1}$.

We first use mirror descent to prove the following covering lemma.

Lemma 132. *Let $d \geq \sqrt{n}$, $2 \leq p \leq q < \infty$, $k := \min(1, d/n)$, $t := \sqrt{(p-1)n} \cdot k^{1/p-1/q}$ and $q^* := q/(q-1)$. Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$, we have*

$$\mathcal{N}\Big(\mathcal{A}(B^d_{S_{q^*}}), \frac{t}{n} B^n_\infty\Big) \leq 2^{O(n)}.$$

Proof. Denote $p^* := p/(p-1)$. Theorem 125 implies $\mathcal{N}(\mathcal{A}(B^d_{S_{q^*}}), k^{1/q^*-1/p^*}\mathcal{A}(B^d_{S_{p^*}})) \leq 2^{O(n)}$, so it suffices to show

$$\mathcal{N}\Big(\mathcal{A}(B^d_{S_{p^*}}), \sqrt{\frac{p-1}{n}} B^n_\infty\Big) \leq 2^{O(n)}.$$

This is a direct consequence of Lemma 128 with norm $\|\cdot\|_{S_{p^*}}$, as the Bregman divergence is $D_\Phi(U,0) = \Phi(U) \leq \frac{1}{2(p^*-1)} = \frac{p-1}{2}$ for $\|U\|_{S_{p^*}} \leq 1$. □

Lemma 132 together with Lemma 46 immediately gives the following weaker measure bound, which we then bootstrap to prove the stronger bound in Theorem 121.

Corollary 133. *Let $d \geq \sqrt{n}$, $2 \leq p \leq q < \infty$ and $k := \min(1, d/n)$. Given symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$, define the convex body*

$$K := \Big\{x \in \mathbb{R}^n : \Big\| \sum_{i=1}^n x_i A_i \Big\|_{S_q} \leq 1\Big\}.$$

Then $\gamma_n(\sqrt{(p-1)n} \cdot k^{1/p-1/q} \cdot K) \geq 2^{-O(n)}$.

Proof of Theorem 121. Let $p_0 := \max(2, \log(2rk))$. For $p \leq p_0$ the result follows directly from Corollary 133, so we may assume $p \geq p_0$. Also note that we may assume $rk \geq 1$ since we can increase smaller values of r without changing the bound on the right side. Remark that $\|A_i\|_{S_{p_0}} \leq r^{1/p_0-1/p}\|A_i\|_{S_p} \leq r^{1/p_0-1/p}$ since the matrices have rank at most r. Corollary 133 then implies that the convex body

$$\sqrt{p_0 n} \cdot k^{1/p_0-1/q} \cdot r^{1/p_0-1/p} \cdot K$$

has Gaussian measure $2^{-O(n)}$. Since $\sqrt{p_0 n} \cdot k^{1/p_0-1/q} \cdot r^{1/p_0-1/p} \lesssim \sqrt{p_0 n} \cdot k^{1/p-1/q}$ by the choice of p_0, it follows that

$$\gamma_n(\sqrt{n \max(1, \log(rk))} \cdot k^{1/p-1/q} \cdot K) \geq 2^{-O(n)},$$

so that Theorem 7 and Corollary 8 yield the partial coloring and full coloring, respectively. The factor $(1/2 + 1/p - 1/q)^{-1}$ comes from the contribution of the exponent of n in the geometric sum, analogous to the second part of Corollary 8. □

7.6 Lower Bound Examples for Matrix Discrepancy

In this section, we give a few examples to illustrate the tightness of our results in Theorem 121 for various regimes of the dimension d and rank r of the input matrices.

7.6.1 Low Dimension Regime of $d = \Theta(\sqrt{n})$

In the regime of $d = \Theta(\sqrt{n})$, we have $k = \min(1, d/n) = \Theta(1/\sqrt{n})$ and $r \leq O(\sqrt{n})$ and our partial coloring bound in Theorem 121 is thus $O(n^{1/2+1/2q-1/2p})$. This bound is tight up to constants due to the following example[5].

Lemma 134 (Example: $d = \sqrt{n}$)**.** *Let $d = \sqrt{n}$ be a power of 2, and $2 \leq p \leq q \leq \infty$. There exist matrices[6] $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$ such that $\| \sum_{i=1}^n x_i A_i\|_{S_q} \gtrsim n^{1/2+1/2q-1/2p}$ for any partial coloring $x \in \{\pm 1\}^n$ with $|\{i : |x_i| = 1\}| \geq n/2$.*

Proof. The idea is to construct an orthogonal basis on $\mathbb{R}^{d\times d}$ with $\|A_i\|_F^2 = d$. Let $H \in \mathbb{R}^{d\times d}$ be the Walsh-Hadamard matrix, and $D_1, \ldots, D_d$ be diagonal matrices with $(D_i)_{j,j} := H_{i,j}$. Let $P_1, \ldots, P_d$ be disjoint permutation matrices, i.e. each P_i permutes the standard orthonormal basis $\{e_1, \ldots, e_d\}$ and each pair P_i, P_j have disjoint non-zero entries. For instance, we may take $(P_i)_{j,k} := 1$ if $j - k \equiv i \mod d$ and 0 otherwise. We then define the n matrices $A_{i+dj} := D_i P_j$ for $i, j \in [d]$. Note that these matrices form an orthogonal basis of $\mathbb{R}^{d\times d}$, so for any partial coloring $x \in \{\pm 1\}^n$ with $|\{i : |x_i| = 1\}| \geq n/2$, we have

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_F^2 = \mathrm{Tr}\left(\left(\sum_{i=1}^n x_i A_i \right)^2 \right) = d \cdot \sum_{i=1}^n x_i^2 \geq dn/2.$$

By Hölder's inequality, this implies that

$$\Big\| \sum_{i=1}^n x_i A_i \Big\|_{S_q} \geq d^{1/q-1/2} \cdot \Big\| \sum_{i=1}^n x_i A_i \Big\|_F \gtrsim n^{1/2+1/2q}.$$

Also note that each matrix A_i has all singular values equal to 1, and therefore $\|A_i\|_{S_p} = d^{1/p} = n^{1/2p}$. Scaling the matrices A_i down by a factor of $n^{1/2p}$ proves the lemma. $\square$

[5] Thanks to Aleksandar Nikolov for suggesting this construction.

[6] These matrices can easily be made symmetric in $\mathbb{R}^{2d\times 2d}$.

7.6.2 Rank-1 Matrices and $d \geq n$

In the regime of $r = 1$ and $d \geq n$, we may assume that $p = 2$. Then the discrepancy bound in Theorem 118 is $O(\sqrt{n})$. This bound is again tight up to a constant factor.

Lemma 135 (Example: $r = 1$ and $d = n$)**.** *Let $2 \leq q \leq \infty$. There exist symmetric rank-1 matrices $A_1, \ldots, A_n \in \mathbb{R}^{n\times n}$ with $\|A_i\|_F \leq 1$ such that any partial coloring $x \in [-1,1]^n$ with $|\{i : |x_i| = 1\}| \geq n/2$ satisfies*

$$\Big\| \sum_{i=1}^{n} x_i A_i \Big\|_{S_q} \geq \Big\| \sum_{i=1}^{n} x_i A_i \Big\|_{\mathrm{op}} \gtrsim \sqrt{n}.$$

Proof. For each $i \in [n-1]$, we define the rank-1 matrices $A_i := \frac{1}{2}(e_i + e_n)(e_i + e_n)^\top$ for $i \in [n]$, where $e_i \in \mathbb{R}^n$ is the unit vector with a single 1 in the ith coordinate and 0 elsewhere, and $A_n = 0$. Note that each $\|A_i\|_F = 1$ by definition. For any partial coloring $x \in [-1,1]^n$ with $|\{i : |x_i| = 1\}| \geq n/2$, we have

$$\sum_{i=1}^{n} x_i A_i = \frac{1}{2} \cdot \begin{pmatrix} x_1 & 0 & \cdots & 0 & x_1 \\ 0 & x_2 & \cdots & 0 & x_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & x_{n-1} & x_{n-1} \\ x_1 & x_2 & \cdots & x_{n-1} & \sum_{i=1}^{n-1} x_i \end{pmatrix}.$$

It then follows that

$$\Big\| \sum_{i=1}^{n} x_i A_i \Big\|_{\mathrm{op}} \geq \Big\| \sum_{i=1}^{n} x_i A_i e_n \Big\|_2 \gtrsim \sqrt{n}.$$

This completes the proof of the lemma. □

As an immediate corollary of Lemma 135, we obtain an $\Omega(\sqrt{n})$ lower bound for matrix Spencer when $d = n$ and all matrices are rank-1.

Corollary 122 (Rank-1 Matrix Spencer Lower Bound)**.** *There exist rank-1 symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{n\times n}$ with $\|A_i\|_{\mathrm{op}} \leq 1$ such that any $x \in \{\pm 1\}^n$ has $\| \sum_{i=1}^n x_i A_i\|_{\mathrm{op}} \gtrsim \sqrt{n}$.*

Another immediate consequence of Lemma 135 is a lower bound of $\Omega(\sqrt{\min(d,n)})$ for Schatten-2 to operator norm discrepancy, which is the generalization of the Komlós problem to matrices. This shows that the Komlós conjecture, which states that the ℓ_2 to ℓ_∞ vector discrepancy is upper bounded by a universal constant, cannot be true for matrices.

Corollary 123 (Lower Bound for Matrix Komlós)**.** *For any d and n, there exist symmetric matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_F \leq 1$ such that any $x \in \{\pm 1\}^n$ has $\|\sum_{i=1}^n x_i A_i\|_{\mathrm{op}} \gtrsim \sqrt{\min(d,n)}$.*

7.7 *An Application of Banaszczyk's Theorem*

We give an alternative simpler proof of the $O(d^{1+1/q-1/p})$ bound for S_p to S_q matrix discrepancy when $d = O(\sqrt{n})$ applying Theorem 12 to a suitable scaling of the operator norm ball:

Corollary 136. *Let $2 \leq p \leq q \leq \infty$. Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_{S_p} \leq 1$, there exists $x \in \{\pm 1\}^n$ such that $\|\sum_{i=1}^n x_i A_i\|_{S_q} \lesssim d^{1+1/q-1/p}$.*

Proof. Note that $\|A_i\|_{S_p} \leq 1$ implies $\|A_i\|_{S_2} \leq d^{1/2-1/p}$. It is well-known that $\gamma_d(4d^{1/2} \cdot B^d_{\mathrm{op}}) \geq 1/2$ (see Theorem 7.3.1 of [167]). Thus, Theorem 12 yields some $x \in \{\pm 1\}^n$ such that $\sum_{i=1}^n x_i A_i \in O(d^{1-1/p}) \cdot B^d_{\mathrm{op}}$. It follows that $\|\sum_{i=1}^n x_i A_i\|_{S_q} \leq O(d^{1+1/q-1/p})$. □

Corollary 137 (Matrix Komlós)**.** *Given matrices $A_1, \ldots, A_n \in \mathbb{R}^{d\times d}$ with $\|A_i\|_F \leq 1$, there exists $x \in \{\pm 1\}^n$ such that $\|\sum_{i=1}^n x_i A_i\|_{S_q} \lesssim \sqrt{\min(d,n)}$, matching the lower bound in Corollary 123.*

Proof. It suffices to take the best between a random coloring, which has discrepancy $O(\sqrt{n})$, and that of Corollary 136. □

Chapter 8

PREFIX AND WEIGHTED DISCREPANCY

This chapter is motivated by discussions with Sander Borst, Daniel Dadush, Haotian Jiang and Lars Rohwedder in 2022.

8.1 Weighted discrepancy

Given a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, its *discrepancy* is a measure of how much balance is achievable when dividing its columns into two groups, defined as $\mathrm{disc}(\boldsymbol{A}) := \min_{\boldsymbol{x} \in \{\pm 1\}^n} \|\boldsymbol{A}\boldsymbol{x}\|_\infty$, where $\|\boldsymbol{z}\|_\infty := \max_{i \in [m]} |z_i|$ is the ℓ_∞ norm. A more robust notion of discrepancy is the *hereditary discrepancy*, given by

$$\mathsf{herdisc}(\boldsymbol{A}) := \max_{\boldsymbol{W} \in \mathrm{diag}\{0,1\}^n} \mathrm{disc}(\boldsymbol{A}\boldsymbol{W}) = \max_{S \subseteq [n]} \mathrm{disc}(\boldsymbol{A}_S),$$

where $\boldsymbol{A}_S \in \mathbb{R}^{m \times S}$ is the matrix A restricted to columns from S.

In this note, we explore the relationship between $\mathsf{herdisc}(\boldsymbol{A})$ and a weighted notion, which we denote $\mathsf{wdisc}(\boldsymbol{A}) := \max_{\boldsymbol{W} \in \mathrm{diag}[0,1]^n} \mathrm{disc}(\boldsymbol{A}\boldsymbol{W})$. Certainly we have $\mathsf{wdisc}(\boldsymbol{A}) \geq \mathsf{herdisc}(\boldsymbol{A})$. A plausible conjecture is the following:

Conjecture 16. *For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ we have $\mathsf{wdisc}(\boldsymbol{A}) < 2 \cdot \mathsf{herdisc}(\boldsymbol{A})$.*

One application of this conjecture is that it would imply that the standard linear programming relaxation of the restricted assignment variant of the max flow time scheduling problem has constant integrality gap [26]. We try to get some intuition of why this conjecture might hold.

We will focus on the case of *totally unimodular* (TU) matrices.

Definition 138. *A matrix $\boldsymbol{A} \in \{-1, 0, 1\}^{m \times n}$ is TU if it satisfies any of the following equivalent properties:*

1. $\mathsf{herdisc}(\boldsymbol{A}) = 1$ *or* $\boldsymbol{A} = \boldsymbol{0}$.

2. $|\det(\boldsymbol{A}_{I\times J})| \leq 1$ *for all* $I \subseteq [m], J \subseteq [n]$ *with* $|I| = |J|$.

3. *For all* $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{Z}^m$ *and* $\boldsymbol{c}, \boldsymbol{m} \in \mathbb{Z}^n$, *the polyhedron* $\{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{c} \leq \boldsymbol{x} \leq \boldsymbol{m}, \boldsymbol{a} \leq \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}\}$ *has only integral vertices.*

TU matrices are ubiquitous in combinatorial optimization, and proving Conjecture 1 holds for them is already enough for the applications.

Conjecture 16 is essentially tight even for TU matrices:

Claim 139. *There exists a TU matrix* $\boldsymbol{A} \in \{0,1\}^{n\times n}$ *so that* $\mathsf{wdisc}(\boldsymbol{A}) = 2 - \frac{2}{n+1}$ *for all odd* n.

Proof. Take $\boldsymbol{A} \in \mathbb{R}^{(2k+1)\times(2k+1)}$ so that $A_{i,j} = 1$ if $i = j$, $i = 2k+1$ or $j = i+1$. For example for $k = 2$ we have the following:

$$\begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

Take weights $w_i = 1$ for odd i and $w_i = 1 - \frac{1}{k+1}$ for even i. Then either two adjacent columns have the same sign and we have weighted discrepancy at least $2 - \frac{1}{k+1} = 2 - \frac{2}{n+1}$ or the signs alternate and the last row gives discrepancy $k + 1 - k(1 - \frac{1}{k+1}) = 2 - \frac{2}{n+1}$. □

We can confirm the conjecture for two classes of TU matrices: formed by consecutive ones and incidence matrices of graphs.

Claim 140. *If* $\boldsymbol{A} \in \{0,1\}^{m\times n}$ *has rows with consecutive ones, then* $\mathsf{wdisc}(\boldsymbol{A}) < 2$.

Proof. Greedily balance all prefixes and write intervals as a difference of two prefixes. □

Theorem 141. *Given an undirected graph $G = ([n], E)$ with $m := |E|$ consider the incidence matrix $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ where $A_{u,e} = 1$ if $u \in e = \{u, v\}$ with $u < v$, $A_{u,e} = -1$ if $u \in e = \{u, v\}$ with $u > v$ and 0 otherwise. Then* $\mathsf{wdisc}(\boldsymbol{A}) \leq 1$.

Proof. Let $w : E \to [0, 1]$ and note that it's equivalent to find an orientation $\gamma : E \to V$ so that $\gamma(e) \in e$ for all $e \in E$ and $\Big|\sum_{v\in e,\gamma(e)=v} w(e) - \sum_{v\in e,\gamma(e)\neq v} w(e)\Big| \leq 1$ for all $v \in V$.

We show the result in fact holds even if the edges $e = \{u, v\}$ are allowed to have different weights on the endpoints $(w(e, u), w(e, v))$, self-loops are allowed and we can have multiple edges between each pair of vertices. First note that self-loops can be dealt with in the end as it amounts to independently balancing numbers in $[-1, 1]$ at each vertex. We proceed by induction on $|V|$. For $|V| = 1$ there can only be self-loops. For $|V| > 1$ consider a vertex v and sort its edges $e_1, \dots, e_d$ according to their weights $w(e_1, v) \geq w(e_2, v) \geq \cdots \geq w(e_d, v)$ with neighbors $u_1, \dots, u_d$. Then since $1 \geq \sum_{i=1}^{d}(-1)^{i+1}w(e_i, v) \geq 0$ in fact any orientation with $\gamma(e_{2i}) \neq \gamma(e_{2i-1})$ will satisfy the discrepancy bound at v. Therefore we may replace the edges e_{2i}, e_{2i-1} by an edge e' between u_{2i} and u_{2i-1} with weights $(w(e_{2i}, u_{2i}), w(e_{2i-1}, u_{2i-1}))$, along with possibly one additional self-loop if d is odd. So by induction we can indeed satisfy the discrepancy bound at all vertices. □

A more difficult class of matrices is that of permutation prefix incidence matrices of graphs. Here instead of a single row for each vertex $v \in V$, we have $d(v)$ many rows, each corresponding to a prefix of the edges incident to v according to some permutation $\pi_v : \{e : v \in e\} \to [d(v)]$. As such, $\boldsymbol{A} \in \mathbb{R}^{(\sum_{v\in V} d(v))\times m} = \mathbb{R}^{2m\times m}$. It can be seen that $\mathsf{herdisc}(\boldsymbol{A}) \leq 1$ so $\boldsymbol{A}$ is still TU. We highlight two cases:

- When $|V| = 2$, this is the *two permutations problem* where given $\pi_1, \pi_2 \in S_m$ we have $2m$ rows corresponding to the prefixes of each permutation;
- When $\pi_v = \pi$ is independent of v, this is the *2-sparse prefix Beck-Fiala problem*.

We show in both of these special cases the conjectured bound of 2 is essentially tight:

Theorem 142. *For $m = 2k + 1$, there exists a matrix $\boldsymbol{A} \in \mathbb{R}^{2m\times m}$ whose rows are prefixes of two permutations so that* $\mathsf{wdisc}(\boldsymbol{A}) = \frac{2k+2}{k+2}$.

Proof. The construction consists of the identity permutation together with the permutation $\pi \in S_{2k+1}$ given by $\pi(2k+1) = 1$, $\pi(i) = 2k+3-i$ if i is even and $\pi(i) = 2k+1-i$ otherwise. For example, for $k = 3$ we have the following two permutations:

$$1\,2\,3\,4\,5\,6\,7$$

$$6\,7\,4\,5\,2\,3\,1$$

Let $\boldsymbol{A}$ denote the associated two permutations matrix and consider the weights given by $w(i) = t+1$ for odd i and $w(i) = t$ for even i with associated diagonal matrix $\boldsymbol{W}$. Eventually we will choose $t := k + 1$. If the signs are alternating $x(i) = (-1)^i$ we get discrepancy $(k + 1)(t + 1) - kt = t + k + 1$ which is $2k + 2$ for $t = k + 1$. So it suffices to show the following:

Claim 143. *If $t \geq k+1 \geq 2$ and $\boldsymbol{x} \in \{\pm 1\}^{2k+1}$ satisfies $x_i = x_{i+1}$ for some i then $\|\boldsymbol{A}\boldsymbol{W}\boldsymbol{x}\|_\infty \geq 2t$.*

For example, the following non-alternating signs yield weighted discrepancy $2t$ at the fourth prefix of each permutation:

$$+1\ -2\ -3\ -4\ +5$$

$$-4\ +5\ -2\ -3\ +1$$

We proceed by induction on k with base cases $k \in \{1, 2\}$, which are easy to check. Suppose that there indeed exist signs $\boldsymbol{x} \in \{\pm 1\}^{2k+1}$ so that $\|\boldsymbol{A}\boldsymbol{W}\boldsymbol{x}\|_\infty < 2t$.

First we show by contradiction that there cannot exist four consecutive indices starting at some even index $2i$ with $x_{2i}w_{2i} + x_{2i+1}w_{2i+1} + x_{2i+2}w_{2i+2} + x_{2i+3}w_{2i+3} = 0$. For otherwise removing such four indices either leads to an instance of the claim for $k-2$, in case their removal leads to non-alternating signs, or else it leads to alternating signs, say starting with $x_1 = +1$. Then $\sum_{j=1}^{2i-1} x_j w_j = t + i$, so that $x_{2i} = -1$ in order to maintain weighted discrepancy less than $2t$. Since the sum is zero, this implies $x_{2i+2} = +1$. Further, $\sum_{j=2i+4}^{2k+1} x_j w_j > 0$ by alternation. This is a prefix of the second permutation, and its following terms are

$x_{2i+2}w_{2i+2}$ and $x_{2i+3}w_{2i+3}$, which cannot both be positive to maintain weighted discrepancy below $2t$, therefore $x_{2i+3} = -1$ and $x_{2i+1} = +1$. But then $\sum_{j=1}^{2i+2} x_j w_j = t+i-t+t+1+t > 2t$, a contradiction.

Thus we may assume no four consecutive indices starting at an even index have weighted signed sum zero. Consider the maximal index i so that $x_i = x_{i+1}$ and assume without loss of generality that $x_{2k+1} = +1$. If i is even, we have $x_i = x_{i+1} = +1$ and $\sum_{j=i+2}^{2k+1} x_j w_j > 0$ by maximality. But then there is a prefix of the second permutation ending at $x_{i+1}w_{i+1}$ with total sum $\sum_{j=i+2}^{2k+1} x_j w_j + x_i w_i + x_{i+1}w_{i+1} > 2t$, so in fact i must be odd and $x_i = x_{i+1} = -1$. If $i = 1$ the prefix $x_1w_1 + x_2w_2$ already violates the discrepancy; if $i = 3$ then either $x_1 = x_2$ and $x_1w_1 + x_2w_2$ is large, or $x_1 \neq x_2$ and the first four indices have discrepancy at least $2t$. Otherwise, $i \geq 5$. By the assumption that no four indices sum to zero, since $x_{i+2} = +1$, we must have $x_{i-1} = -1$. Remark that no four consecutive indices can have the same sign, so that $x_{i-2} = +1$. Again by the assumption that no four sum to zero, $x_{i-3} = -1$. But then the second permutation has a prefix of value $\sum_{j=i+2}^{2k+1} x_j w_j - (4t+2) = \frac{2k+1-(i+2)}{2} + t + 1 - (4t+2) \leq k - 3 - 3t - 1 < -2t$, a contradiction. □

This also yields a construction for 2-sparse prefix Beck-Fiala:

Theorem 144. *There exists an instance $\boldsymbol{A}$ of 2-sparse prefix Beck-Fiala with $2k+3$ vertices and $4k+2$ edges with* $\mathsf{wdisc}(\boldsymbol{A}) \geq \frac{2k+2}{k+2}$.

Proof. The construction is a complete bipartite graph $K_{2,2k+1}$ where the two vertices on the left side correspond to each of the two permutations on $2k+1$ elements from Theorem 142 with the same alternating weights of $k+1$ and $k+2$. We construct a single global permutation where the edges are listed in the order of the first permutation followed by the second. The only observation needed is that for each vertex on the right side, if both edges are oriented the same way then the discrepancy is already at least $2(k+1) = 2k+2$. Otherwise, all pairs of edges connected to vertices in the right side are in opposite orientations, so that a lower bound for the two permutations translates to a lower bound for the prefix discrepancy in the incidence matrix of the graph. □

Chapter 9

OPTIMAL ONLINE DISCREPANCY MINIMIZATION

This chapter is based on joint work with Janardhan Kulkarni and Thomas Rothvoss [90].

9.1 Introduction

We study *online* vector balancing problems, first considered by Spencer in late 70's [156]. We receive vectors $v_1, \ldots, v_T \in \mathbb{R}^n$, which are bounded in some norm, one at a time, and we have to decide the sign $x_i \in \{-1, 1\}$ *irrevocably* after learning the vector v_i. The goal is to keep the signed sum $\sum_{i=1}^{T} x_i v_i$ small in some norm; a natural variant asks that all prefixes $\sum_{i=1}^{t} x_i v_i$ are small for all $t \in [T]$. This vector balancing formulation captures several classic problems in discrepancy theory, where the norm to be balanced is the maximum absolute value of any coordinate, also known as the ℓ_∞ norm.

If $\|v_i\|_\infty \leq 1$, uniformly random signs achieve a $O(\sqrt{T \log n})$ ℓ_∞ bound, and there are several methods to make this a deterministic online algorithm, see for example the excellent book of Chazelle [45]. Unfortunately, the random coloring is tight in its dependency on $\sqrt{T}$ for $n \geq 2$ [156, 158]. Indeed, an *adaptive adversary* can simply choose a vector $v_t \in [-1, 1]^n$ that is orthogonal to the current position $\sum_{i=1}^{t-1} x_i v_i$ and satisfies $\|v_t\|_2^2 \geq n - 1$, then $\| \sum_{i=1}^{t} x_i v_i \|_2 \geq \sqrt{(n-1)T}$ and, in particular, the discrepancy is $\Omega(\sqrt{T})$.

Much of the focus of subsequent efforts has been to improve the dependence on the $\sqrt{T}$ term by restricting the power of the adversary. One natural choice is to consider the stochastic setting. Here vectors v_i are sampled independently from a distribution $\mathbf{p}$ that is known to the online algorithm. When $\mathbf{p}$ is a uniform distribution on all $\{-1, 1\}^n$ vectors, Bansal and Spencer [27] showed that one can get $O(\sqrt{n})$ discrepancy for the ℓ_∞-norm, or $O(\sqrt{n} \log T)$ for all prefixes up to time T. Motivated by the applications of online discrepancy minimization techniques to online envy minimization problems [78], Bansal, Jiang, Singla and Sinha [24] considered general distributions $\mathbf{p}$ supported on $[-1,1]^n$. For this problem, they achieved an ℓ_∞-discrepancy of $O(n^2 \log(nT))$, and this was improved by Bansal, Jiang, Meka, Singla, and Sinha [22] who showed a bound of $O(\sqrt{n}\log^4(nT))$. The work of Aru, Narayanan, Scott and Venkatesan [8] achieved a bound of $O_n(\sqrt{\log T})$, where the dependence on n is super exponential. One important point to note in all of these results is that they substantially improve the dependence on $\Omega(\sqrt{T})$ to logarithmic factors.

Despite these impressive results, until very recently, little progress was made on the online vector balancing problem against *oblivious* adversaries – the most common setting considered in the online algorithms literature. Here the adversary fixes an arbitrary set of vectors $v_1, \ldots, v_T \in \mathbb{R}^n$ in advance, and the online algorithm can use randomized strategies. For the special case of *edge orientation*, where the vectors correspond to columns of the incidence matrix of a graph, a simple random labeling argument due to Kalai [120] achieves a discrepancy bound of $O(\log T)$ with high probability.

In an elegant result, Alweiss, Liu and Sawhney [5] showed that a very simple *self-balancing random walk* can find random signs so that with high probability all prefixes $\sum_{i=1}^t x_i v_i$ are $O(\sqrt{\log(nT)})$-*subgaussian*. Here a random vector $X \in \mathbb{R}^n$ is called c-*subgaussian* if for any $w \in S^{n-1}$ one has $\mathbb{E}[\exp(\langle X, w\rangle^2/c^2)] \leq 2$. In particular, all prefix sums $\|\sum_{i=1}^t x_i v_i\|_\infty$ are $O(\log(nT))$ with high probability against any oblivious adversary.

9.1.1 Our contributions

Our main contribution is:

Theorem 145. *There is an online algorithm that against any oblivious adversary and for any sequence of vectors $v_1, \ldots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$, arriving one at a time, decides random signs $x_1, \ldots, x_T \in \{-1, 1\}$ so that for every $t \in [T]$, the prefix sum $\sum_{i=1}^t x_i v_i$ is 10-subgaussian.*

The algorithm does not depend on T so one may also take an infinite sequence of vectors.

By using the machinery of Talagrand's majorizing measures theorem, we may recover Banaszczyk's theorem in the online setting.

Theorem 146. *Given a symmetric convex body $K \subseteq \mathbb{R}^n$, there is an online algorithm that against any oblivious adversary and for any sequence of vectors $v_1, \dots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \le 1$, arriving one at a time, decides random signs $x_1, \dots, x_T \in \{-1, 1\}$ so that each of the following hold with probability at least $1/2$:*

(a) $\sum_{i=1}^T x_i v_i \in O(1) \cdot K$ *under the assumption* $\gamma_n(K) \ge \frac{1}{2}$.

(b) $\sum_{i=1}^t x_i v_i \in O(1) \cdot K$ *for all* $t \in [T]$ *under the assumption* $\gamma_n(K) \ge 1 - \frac{1}{2T}$.

For ℓ_p discrepancy minimization, we obtain the following corollary:

Corollary 147. *There is an online algorithm that against any oblivious adversary and for any sequence of vectors $v_1, \dots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \le 1$, arriving one at a time, decides random signs $x_1, \dots, x_T \in \{-1, 1\}$ so that each of the following hold with probability at least $1 - \delta$ for any $\delta \in (0, \frac{1}{2}]$ and any $p \ge 2$:*

(a) $\|\sum_{i=1}^T x_i v_i\|_p \lesssim \sqrt{p}\min(n, T)^{1/p} + \sqrt{\log(1/\delta)}$;

(b) $\max_{t \in [T]} \|\sum_{i=1}^t x_i v_i\|_p \lesssim \sqrt{p}\min(n, T)^{1/p} + \sqrt{\log T} + \sqrt{\log(1/\delta)}$.

Furthermore,

(c) $\|\sum_{i=1}^T x_i v_i\|_\infty \lesssim \sqrt{\log \min(n, T)} + \sqrt{\log(1/\delta)}$;

(d) $\max_{t \in [T]} \|\sum_{i=1}^t x_i v_i\|_\infty \lesssim \sqrt{\log T} + \sqrt{\log(1/\delta)}$.

Our bounds match the best known upper bounds in the offline setting where all the vectors are known in advance. We show Corollary 147(d) is tight in the oblivious setting even when $n = 2$:

Theorem 148. *For any $n \ge 2$, there is a strategy for an oblivious adversary that yields a sequence of unit vectors $v_1, \dots, v_T \in \mathbb{R}^n$ so that for any online algorithm, with probability at least $1 - 2^{-T^{\Omega(1)}}$, one has $\max_{t \in [T]} \|\sum_{i=1}^t x_i v_i\|_\infty \gtrsim \sqrt{\log T}$.*

This improves upon the $\Omega\Big(\sqrt{\frac{\log T}{\log\log T}}\Big)$ lower bound of [24]. For the ℓ_2 norm, Corollary 147(a) is tight for any n, T, since an orthonormal basis followed by zeros does achieve a lower bound of $\sqrt{\min(n,T)}$. If the orthonormal basis is instead followed by the construction in 148, we get a matching lower bound construction for Corollary 147(b). On the other hand, it is not known whether there exists a lower bound better than constant for Corollary 147(c).

Finally, our framework also provides improved bounds for the edge orientation problem.

Corollary 149 (Online edge orientation)**.** *There exists an online algorithm that for any set of n vertices and any sequence of edges, arriving one at a time, decides orientations so that at every vertex, the absolute difference between indegree and outdegree always remains bounded by $O(\sqrt{\log T})$ with high probability.*

9.1.2 An Overview of our Techniques

Suppose we play against an oblivious adversary that has a predetermined sequence of vectors $v_1, \ldots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$ for all $i \in [T]$ that are revealed to us one vector at a time and we need to determine random signs $x_i \in \{-1, 1\}$. After some discretization one may think of this game as a balancing problem on a rooted tree $\mathcal{T} = (V, E)$ where edges $e \in E$ are labelled with vectors v_e of length at most one. The adversary chooses a path from the root to a leaf which is revealed one edge at a time. After learning the next edge on the path we must give it a random sign $x_e \in \{-1, 1\}$ to keep the sum $\sum_{e \in P} x_e v_e$ subgaussian where P is the path chosen so far by the adversary. Here the tree $\mathcal{T}$ has depth T but each interior node will have an outgoing edge for each vector in a fine enough ε-net and so degrees in $\mathcal{T}$ are exponential in n and T.

The next observation is that rather than only drawing signs for the selected path we can draw all signs $x \in \{-1, 1\}^E$ with the goal of keeping $\sum_{e \in P} x_e v_e$ $O(1)$-subgaussian for *all* root-node paths P. This is not actually a restriction since the adversary could pick any such path anyway. Note that indeed this argument assumes an oblivious adversary. In order to find the distribution we want to make use of Banaszczyk's Theorem 11.

Banaszczyk's construction for the body $K * u$ is quite intuitive: call a line $x + \mathbb{R}u$ *long* if its intersection with K has length at least $2\|u\|_2$. Then intersect $(K + u) \cup (K - u)$ with all long lines. However, the proof of the inequality $\gamma_n(K * u) \geq \gamma_n(K)$ is quite skillful. The main application in Banaszczyk [14] was to prove that for any vectors $v_1, \ldots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$ and any convex body $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq 1/2$, there are signs $x \in \{-1, 1\}^T$ with $\sum_{i=1}^T x_i v_i \in 5K$. In a later work, Banaszczyk applied his Theorem 11 also to the prefix setting (Theorem 13). One can think of this result as balancing all root-node paths on a path graph into a body K. It is also interesting to note that the construction behind Theorem 13 is inherently non-online — the sign for v_i may depend on all other vectors $v_1, \ldots, v_{i-1}$ and $v_{i+1}, \ldots, v_T$. Our first step is to generalize Theorem 13 to trees $\mathcal{T} = (V, E)$ and prove that for any convex body K with $\gamma_n(K) \geq 1 - \frac{1}{2|E|}$ we can find signs $x \in \{-1, 1\}^E$ so that $\sum_{e \in P} x_e v_e \in \alpha K$ for any root-node path P. However, there are two problems with that approach:

(i) the tree $\mathcal{T}$ obtained in the reduction has exponential size and K would need to be huge to satisfy $\gamma_n(K) \geq 1 - \frac{1}{2|E|}$;

(ii) a straightforward application of Banaszczyk's framework will only produce a single vector of signs $x \in \{-1, 1\}^E$ rather than a distribution.

Interestingly, both problems can be solved with the same *cloning trick*: we replace each edge e with a sequence of N edges labelled with orthogonal copies of v_e, say $v_e^{(1)}, \ldots, v_e^{(N)}$. Then the signs $x_e^{(\ell)} \in \{-1, 1\}$ for the blown-up tree can be turned into a distribution over signs $(x_e^{(\ell)})_{e \in E}$ for the original tree by drawing $\ell \sim \{1, \ldots, N\}$ uniformly. This solves issue (ii). Considering (i), our trick makes the tree even larger by a factor N, worsening the issue! However, choosing

$$K = \left\{ (y_e^{(\ell)})_{e \in E, \ell \in [N]} \mid (y_e^{\ell})_{e \in E} \text{ is } O(1)\text{-subgaussian when } \ell \sim [N] \right\}$$

we can prove that the Gaussian measure of K increases fast enough as N grows to compensate for the blowup of the tree.

9.2 Preliminaries

First, we review a few facts for later.

9.2.1 Geometry

A *convex body* is a set $K \subseteq \mathbb{R}^n$ that is convex, compact (closed and bounded) and full-dimensional. Let $B_2^n := \{x \in \mathbb{R}^n \mid \|x\|_2 \leq 1\}$ be the *Euclidean ball* and let $S^{n-1} := \{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$ be the *sphere*. A set $W \subseteq S^{n-1}$ is called an *ε-net* if for all $x \in S^{n-1}$, there is a $y \in W$ with $\|x - y\|_2 \leq \varepsilon$.

Lemma 150. *For any $0 < \varepsilon \leq 1$, there is an ε-net $W \subseteq S^{n-1}$ of size $|W| \leq (\frac{3}{\varepsilon})^n$.*

Proof. Pick any maximal set of points $W \subseteq S^{n-1}$ that have $\|\cdot\|_2$-distance at least ε to each other. Then W is an ε-net. Moreover the balls $x + \frac{\varepsilon}{2} B_2^n$ are disjoint for $x \in W$ and contained in $(1 + \frac{\varepsilon}{2}) B_2^n$. Hence

$$|W| \leq \frac{\mathrm{Vol}_n((1 + \frac{\varepsilon}{2}) \cdot B_2^n)}{\mathrm{Vol}_n(\frac{\varepsilon}{2} \cdot B_2^n)} = \left(\frac{1 + \frac{\varepsilon}{2}}{\frac{\varepsilon}{2}}\right)^n \leq \left(\frac{3}{\varepsilon}\right)^n. \qquad \square$$

Lemma 151. *Let $W \subseteq S^{n-1}$ be an ε-net for $0 < \varepsilon < 1$. Then for any $w_0 \in S^{n-1}$ there is a coefficient vector $\lambda \in \mathbb{R}^W_{\geq 0}$ with $w_0 = \sum_{w \in W} \lambda_w w$ and $\sum_{w \in W} \lambda_w \leq \frac{1}{1-\varepsilon}$.*

Proof. Consider $L(u) := \min\{\|\lambda\|_1 : u = \sum_{w \in W} \lambda_w w \text{ and } \lambda_w \geq 0 \ \forall w \in W\}$. It is not entirely obvious, but $L(u)$ is well-defined for each $u \in S^{n-1}$. To see this, note that $L(u)$ is the value of a linear program and if that program was infeasible then by Farkas Lemma (see e.g. Schrijver [154]) there is a $c \in S^{n-1}$ with $c^T u > 0$ and $c^T w < 0$ for all $w \in W$. Then $\|c - w\|_2 \geq 1$ for all $w \in W$ which is a contradiction. Now, fix a $w^* \in S^{n-1}$ maximizing the value $L(w^*)$; such a point must exist by compactness of S^{n-1}. Let $w \in W$ with $\|w^* - w\|_2 \leq$

ε. Then writing $w^* = w + (w^* - w)$ we have

$$L(w^*) \leq \underbrace{L(w)}_{\leq 1} + \underbrace{\|w^* - w\|_2}_{\leq \varepsilon} \cdot L\Big(\frac{w^* - w}{\|w^* - w\|_2}\Big) \leq 1 + \varepsilon L(w^*).$$

Rearranging gives $L(w_0) \leq L(w^*) \leq \frac{1}{1-\varepsilon}$. □

9.2.2 Probability

We need to formalize what we mean that a random variable has Gaussian-type tails. There are several equivalent ways to do so. We recommend the excellent textbook by Vershynin [167] for background.

Proposition 152 ([167])**.** *Let X be a random variable. The following properties are equivalent; the parameters $C_i > 0$ appearing in these properties differ from each other by at most an absolute constant factor.*

(*i*) *One has* $\Pr[|X| \geq t] \leq 2\exp(-t^2/C_1^2)$ *for all* $t \geq 0$*.*

(*ii*) *One has* $\mathbb{E}[\exp(X^2/C_2^2)] \leq 2$*.*

If additionally $\mathbb{E}[X] = 0$*, then* $(i) + (ii)$ *are also equivalent to*

(*iii*) *One has* $\exp(\lambda X) \leq \exp(C_3^2\lambda^2)$ *for all* $\lambda \in \mathbb{R}$*.*

While Dadush et al [51] and Bansal et al [19] made use of (iii), for our purpose, it will be most useful to work with (ii). Hence for a real-valued random variable X, we define the *subgaussian norm* as

$$\|X\|_{\psi_2} := \inf\ \big\{t > 0 : \mathbb{E}[\exp(X^2/t^2)] \leq 2\big\}.$$

It can be proven that $\|\cdot\|_{\psi_2}$ is a norm on the space of jointly distributed random variables; indeed it is the *Orlicz norm* associated with the (convex and nondecreasing) function $z \mapsto e^{z^2} - 1$. In particular $\|X + Y\|_{\psi_2} \leq \|X\|_{\psi_2} + \|Y\|_{\psi_2}$ for any random variables X, Y — even if they are not independent.

For the sake of completeness and in order to be able to give explicit constants, we show the directions $(ii) \Rightarrow (i)+(iii)$ in Proposition 152. For example, the following lemma shows that any 10-subgaussian random variable X satisfies $\mathbb{E}[\exp(\lambda X)] \leq \exp(20\lambda^2)$. Therefore Theorem 145 also yields the subgaussianity bound given in the main result of [19].

Lemma 153. *If* $\|X\|_{\psi_2} \leq 1$ *then* $\Pr[|X| \geq t] \leq 2e^{-t^2}$ *for all* $t \geq 0$*. If moreover* $\mathbb{E}[X] = 0$*, then also* $\mathbb{E}[\exp(\lambda X)] \leq \exp(0.4\lambda^2)$ *for all* $\lambda \in \mathbb{R}$*.*

Proof. The claim that $\Pr[|X| \geq t] \leq 2e^{-t^2}$ follows directly from Markov's inequality. For the second claim, if $0.56\lambda^2 < 1$ we can use the inequality $e^z \leq z + e^{0.56z^2}$ (valid for all $z \in \mathbb{R}$) to obtain

$$\mathbb{E}[\exp(\lambda X)] \leq \mathbb{E}[\exp(X^2)^{0.56\lambda^2}] \overset{\text{concavity}}{\leq} \underbrace{\mathbb{E}[\exp(X^2)]}_{\leq 2 \text{ as } \|X\|_{\psi_2} \leq 1}{}^{0.56\lambda^2} \leq \exp(0.4\lambda^2).$$

Else, $\lambda^2 \geq \frac{1}{0.56}$ and we will show that $e^z \leq z + \frac{1}{2}\exp(0.4\lambda^2 + \frac{z^2}{\lambda^2})$ for all $z \in \mathbb{R}$ to obtain

$$\mathbb{E}[\exp(\lambda X)] \leq \tfrac{1}{2}\exp(0.4\lambda^2)\,\mathbb{E}[\exp(X^2)] \leq \exp(0.4\lambda^2).$$

For $\lambda^2 = \frac{1}{0.56}$, the claim follows from $e^z \leq z + e^{0.56z^2}$; the only other possible minimizer λ_* of the right side occurs when $0.4\lambda_*^2 = \frac{z^2}{\lambda_*^2}$, so $\lambda_*^4 = z^2/0.4 \geq \frac{1}{0.56^2}$. And indeed, we do have the inequality $e^z \leq z + \frac{1}{2}\exp(2\sqrt{0.4}|z|)$ for all $|z| \geq \frac{\sqrt{0.4}}{0.56}$. $\square$

Next, we extend the concept of a subgaussian norm from random variables to random vectors. For a random vector X taking values in $\mathbb{R}^n$, we denote

$$\|X\|_{\psi_2,\infty} := \sup_{w \in S^{n-1}} \|\langle X, w\rangle\|_{\psi_2}.$$

Again, $\|\cdot\|_{\psi_2,\infty}$ is a norm on the space of random vectors in $\mathbb{R}^n$. For simplicity of notation, we sometimes write that X is *t-subgaussian* if $\|X\|_{\psi_2,\infty} \leq t$. Of course, a standard Gaussian $g \sim N(0, I_n)$ is also $O(1)$-subgaussian, but curiously not with constant 1.

Lemma 154. *For* $\lambda < \frac{1}{2}$ *one has* $\mathbb{E}_{g\sim N(0,1)}[\exp(\lambda \cdot g^2)] = \frac{1}{\sqrt{1-2\lambda}}$*. In particular, the standard Gaussian* $N(\mathbf{0}, I_n)$ *is* $\sqrt{\frac{8}{3}}$*-subgaussian.*

We will also need the following upper bound on moments of sums of independent random variables:

Lemma 155 (Rosenthal's inequality [144])**.** *Let $p \geq 2$ and let $X_1, \dots, X_N$ be mean zero independent random variables with finite p-th moment $\mathbb{E}[|X_i|^p]$. Then*

$$\mathbb{E}[|X_1 + \cdots + X_N|^p]^{1/p} \leq 2^p \cdot \max\Big\{\Big(\sum_{i=1}^N \mathbb{E}[|X_i|^p]\Big)^{1/p}, \Big(\sum_{i=1}^N \mathbb{E}[X_i^2]\Big)^{1/2}\Big\}.$$

We also mention that the constant 2^p can be improved to $\Theta(p/\log p)$ [92], though it will not be needed in our application.

Finally, we will use the tail bound form of Talagrand's comparison inequality. See once again the textbook of Vershynin [167] (Chapter 8.6).

Lemma 156. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body and $X \in \mathbb{R}^n$ a random vector that is $O(1)$-subgaussian. Then with probability at least $1 - \delta$ for any $\delta \in (0, \frac{1}{2}]$,*

$$\|X\|_K \lesssim \mathop{\mathbb{E}}_{g \sim N(\mathbf{0}, I_n)}[\|g\|_K] + \sqrt{\log(1/\delta)} \cdot \frac{1}{\mathrm{inradius}(K)},$$

where $\mathrm{inradius}(K)$ is the largest $r > 0$ so that $rB_2^n \subseteq K$.

9.3 Generalizing Banaszczyk's prefix balancing argument to trees

We begin by generalizing Banaszczyk's Theorem 13 to trees. A similar statement may also be found (with vectors assigned to vertices) as Theorem 4.1. in [23].

Theorem 157. *There exists a constant $\alpha < 5$ such that the following holds. Let $\mathcal{T} = (V, E)$ be a tree with a distinguished root $r \in V$ and $|E| \geq 1$, where each edge $e \in E$ is assigned a vector $v_e \in \mathbb{R}^n$ with $\|v_e\|_2 \leq 1$. Let $K \subseteq \mathbb{R}^n$ be a convex body with $\gamma_n(K) \geq 1 - \frac{1}{2|E|}$. Then there are signs $x \in \{-1,1\}^E$ so that*

$$\sum_{e \in P_i} x_e v_e \in \alpha K \quad \forall i \in V$$

where $P_i \subseteq E$ are the edges on the path from the root to i.

Proof. Let $\beta > 1/5$ be the constant from Theorem 11. We write $C_i \subseteq V$ as all the *children* of i where $C_i = \emptyset$ for leaves and we write $D_i \subseteq V$ as the *descendants* of i (including i itself). In particular $\{i\} \cup C_i \subseteq D_i$. We construct a sequence of target bodies by setting

$$K_i := \Big(\bigcap_{j \in C_i} (K_j * \beta v_{\{i,j\}}) \Big) \cap K.$$

Then for any leaf $i \in V$ one simply has $K_i = K$. We prove that the constructed bodies are still large:

Claim I. *For all $i \in V$ one has $\gamma_n(K_i) \geq 1 - \frac{|D_i|}{2|E|}$.*

Proof of Claim I. We prove the claim by induction. The claim is true for any leaf $i \in V$ because in that case $|D_i| = 1$ and $\gamma_n(K_i) = \gamma_n(K) \geq 1 - \frac{1}{2|E|}$. Now consider any interior node i and assume the claim is true for all its children. In particular for each $j \in C_i$ we have $\gamma_n(K_j) \geq 1 - \frac{|D_j|}{2|E|} \geq \frac{1}{2}$ by induction and the fact for any non-root node $|D_j| \leq |E|$. Hence

$$\gamma_n(K_j * \beta v_{\{i,j\}}) \overset{\text{Thm 11}}{\geq} \gamma_n(K_j) \overset{\text{induction}}{\geq} 1 - \frac{|D_j|}{2|E|}.$$

Then

$$\begin{aligned} \gamma_n(K_i) \quad &= \quad \gamma_n\Big(\Big(\bigcap_{j \in C_i} (K_j * \beta v_{\{i,j\}})\Big) \cap K\Big) \\ &\overset{\text{union bound}}{\geq} \quad 1 - \sum_{j \in C_i} \underbrace{\gamma_n(\mathbb{R}^n \setminus (K_j * \beta v_{\{i,j\}}))}_{\leq \frac{|D_j|}{2|E|}} - \underbrace{\gamma_n(\mathbb{R}^n \setminus K)}_{\leq \frac{1}{2|E|}} \\ &\geq \quad 1 - \frac{1}{2|E|} \underbrace{\Big(1 + \sum_{j \in C_i} |D_j|\Big)}_{=|D_i|} = 1 - \frac{|D_i|}{2|E|}. \qquad \square \end{aligned}$$

Claim II. *There are signs $x \in \{-1,1\}^E$ so that $\sum_{e \in P_i} x_e v_e \in \frac{1}{\beta} K_i$ for all $i \in V \setminus \{r\}$.*

Proof of Claim II. We construct the signs in increasing distance to the root. If i is a child of the root itself, then the claim is true as $\gamma_n(K_i) \geq \frac{1}{2}$ and $\|v_{\{r,i\}}\|_2 \leq 1$. Suppose for some node $i \in V \setminus \{r\}$ we have determined all the signs $x_e \in \{-1,1\}$ with $e \in P_i$ and in particular $a := \beta \sum_{e \in P_i} x_e v_e \in K_i$. Consider any child $j \in C_i$. Then

$$a \in K_i \overset{\text{Def } K_i}{\subseteq} K_j * \beta v_{\{i,j\}} \overset{\text{Thm 11}}{\subseteq} (K_j + \beta v_{\{i,j\}}) \cup (K_j - \beta v_{\{i,j\}})$$

Then we may pick a sign $x_{\{i,j\}} \in \{-1, 1\}$ so that $a + \beta x_{\{i,j\}} v_{\{i,j\}} \in K_j$. □

Then the overall claim follows from the fact that $K_i \subseteq K$ for all $i \in V \setminus \{r\}$ and for the root itself one trivially has $\mathbf{0} \in K$. □

9.4 The body of subgaussian distributions

In this section, we introduce the target body K that we will use and we will prove that its Gaussian measure is close enough to 1 to make the argument work. For $n, N \in \mathbb{N}$ and a small $\delta > 0$, we define

$$K := \Big\{(y^{(1)}, \dots, y^{(N)}) \in \mathbb{R}^{Nn} \mid \|Y\|_{\psi_2,\infty} \leq 2 + \delta \text{ where } Y \sim \{y^{(1)}, \dots, y^{(N)}\}\Big\}. \tag{9.1}$$

Intuitively, the vectors in K consist of N many blocks of dimension n with the property that a uniform random block generates a subgaussian random vector. Since $\|\cdot\|_{\psi_2,\infty}$ is a norm, K is a symmetric convex body. The main result for this section will be that K has a large Gaussian measure if N is large enough.

Proposition 158. *For any $\delta > 0$, there is a constant $C_\delta > 0$ so that for all $n, N \in \mathbb{N}$ one has $\gamma_{Nn}(K) \geq 1 - \frac{C_\delta^n}{N^{1+\delta}}$.*

We first show how to control the deviation in a single direction w. Note that random variables for the form $X := \exp(g^2)$ with $g \sim N(0, \sigma^2)$ and $\sigma > 0$ have *heavy tails*, which means they do not possess finite exponential moments (i.e. $\mathbb{E}[e^{\lambda X}] = \infty$ for all $\lambda > 0$). That implies in particular that standard Chernoff bounds cannot be used to derive concentration.

Lemma 159. *For any $C > 4$ and $\delta' \in \Big(0, \frac{C}{4} - 1\Big)$, there is a $C' > 0$ so that for any unit vector $w \in S^{n-1}$, the set*

$$K_w := \Big\{(y^{(1)}, \dots, y^{(N)}) \in \mathbb{R}^{Nn} \mid \mathop{\mathbb{E}}_{\ell \sim [N]}\big[\exp\big(\tfrac{1}{C}\left\langle w, y^{(\ell)}\right\rangle^2\big)\big] \leq 2\Big\}$$

satisfies $\gamma_{Nn}(K_w) \geq 1 - \frac{C'}{N^{C/4-\delta'}}$.

Proof. Since $y^{(\ell)} \sim N(\mathbf{0}, I_n)$, the inner product satisfies $\langle w, y^{(\ell)} \rangle \sim N(0,1)$. We define $Y_\ell := \exp\left(\frac{1}{C}\left\langle w, y^{(\ell)}\right\rangle^2\right)$ with $\mu_C := \mathbb{E}[Y_\ell] = \left(1 - \frac{2}{C}\right)^{-1/2}$ by Lemma 155 and the centered random variable $X_\ell := Y_\ell - \mu_C$, so that for any $p > 2$ we have by Markov

$$\begin{aligned}\Pr\left[Y_1 + \cdots + Y_N > 2N\right] &= \Pr\left[X_1 + \cdots + X_N > (2 - \mu_C) \cdot N\right] \\ &\le \Pr\left[|X_1 + \cdots + X_N|^p > (2 - \mu_C)^p \cdot N^p\right] \\ &\le \frac{\mathbb{E}[|X_1 + \cdots + X_N|^p]}{(2 - \mu_C)^p N^p}.\end{aligned}$$

Since X_ℓ are mean zero independent random variables with p-th moment

$$\mathbb{E}[|X_\ell|^p] = E[|Y_\ell - \mu_C|^p] \le \mathbb{E}[Y_\ell^p] + \mu_C^p = \mu_{C/p} + \mu_C^p,$$

which is a finite constant M_p for $p < C/2$, we may apply Lemma 155 to obtain

$$\mathbb{E}[|X_1 + \cdots + X_N|^p] \le \left(2^p \cdot \max\{M_p \cdot N^{1/p}, M_2 \cdot N^{1/2}\}\right)^p \le C_p N^{p/2}.$$

For $p := C/2 - 2\delta'$, we conclude $\Pr\left[Y_1 + \cdots + Y_N > 2N\right] \le \frac{C_p}{(2-\mu_C)^p N^{p/2}} = \frac{C'}{N^{C/4-\delta'}}$ for some constant C' depending only on C and δ'. $\square$

Remark 6. *Lemma 159 is tight in the sense that one cannot hope for a lower bound $\gamma_{Nn}(K_w) \ge 1 - O(1/N)$ for any $C < 4$. Indeed, for $Y_\ell := \exp\left(\frac{1}{C}\left\langle w, y^{(\ell)}\right\rangle^2\right)$, one has*

$$\Pr\left[Y_1 + \cdots + Y_N > 2N\right] \ge \Pr\left[\max_{\ell \in [N]} Y_\ell > 2N\right] = 1 - \left(\Pr_{g \sim N(0,1)}\left[g \le \sqrt{C \log(2N)}\right]\right)^N,$$

which is $\Omega\left(\frac{1}{N^{C/2-1}\sqrt{\log(2N)}}\right)$ for any $C \ge 2$ by using standard Gaussian tail bounds.

Remark 7. *In fact, the lower bound above is tight up to constants for $C > 8/3$ (which ensures $2 > \mu_C$). This is a consequence of Theorem 3.1.6 in [35] (see Equation 2.1.6 for context). The proof follows from a truncation of the random variables and a careful estimation of the resulting moment generating function. Thus there is a stronger version of Lemma 159 which also works for $C = 4$, although it does not seem to lead to an improvement of Proposition 158.*

Obviously one cannot combine Lemma 159 with a union bound over the infinitely many vectors $w \in S^{n-1}$, but it is a standard argument that in such cases it suffices to control the deviation in directions of an ε-net, see again the textbook of Artstein-Avidan, Giannopoulos and Milman [7].

Lemma 160. *Let $W \subseteq S^{n-1}$ be an ε-net for $0 < \varepsilon < 1$ and let $X \in \mathbb{R}^n$ be a random vector so that $\| \langle X, w \rangle \|_{\psi_2} \leq 1$ for all $w \in W$. Then $\|X\|_{\psi_2,\infty} \leq \frac{1}{1-\varepsilon}$.*

Proof. Fix $w \in S^{n-1}$. By Lemma 151 there are coefficients $\lambda \in \mathbb{R}^W_{\geq 0}$ with $w = \sum_{u \in W} \lambda_u u$ and $\|\lambda\|_1 \leq \frac{1}{1-\varepsilon}$. Then as $\| \cdot \|_{\psi_2}$ is a norm we have

$$\| \langle X, w \rangle \|_{\psi_2} = \Big\| \Big\langle X, \sum_{u \in W} \lambda_u u \Big\rangle \Big\|_{\psi_2} \leq \sum_{u \in W} \lambda_u \underbrace{\| \langle X, u \rangle \|_{\psi_2}}_{\leq 1} \leq \frac{1}{1-\varepsilon}. \qquad \square$$

Now we are ready to prove Proposition 158:

Proof of Prop. 158. Set $C := (2+\delta)^2 - \delta^2/2$, $\varepsilon := 1 - \frac{\sqrt{C}}{2+\delta}$ and let W be an ε-net of S^{n-1} of size at most $(\frac{3}{\varepsilon})^n$ (see Lemma 150). Applying Lemma 159 with C and $\delta' := \delta^2/8$ (chosen so that $C/4 - \delta' = 1 + \delta$) and taking the union bound over all $w \in W$, it follows that for $y^{(\ell)} \sim N(\mathbf{0}, I_n)$ and $Y \sim \{y^{(1)}, \dots, y^{(N)}\}$ one has

$$\Pr[\|\langle Y, w \rangle\|_{\psi_2} \leq \sqrt{C} \; \forall w \in W] \geq 1 - \left(\frac{3}{\varepsilon}\right)^n \cdot \frac{C'}{N^{1+\delta}}$$

for some $C' > 0$. If this event happens, then by Lemma 160 also $\|Y\|_{\psi_2,\infty} \leq 2 + \delta$. $\square$

9.5 *Subgaussianity over trees*

Recall that in Theorem 157 we have proven that we can find a *single coloring* so that all vectors on any path in a tree are balanced into a (large enough) convex set K. We can use the same argument to find a *distribution* that keeps all those vector sums subgaussian.

Theorem 161. *There exists a constant $\gamma < 10$ such that the following holds. Let $\mathcal{T} = (V, E)$ be a tree with a distinguished root, where each edge $e \in E$ is assigned a vector $v_e \in \mathbb{R}^n$ with $\|v_e\|_2 \leq 1$. Then there is a distribution $\mathcal{D}$ over $\{-1,1\}^E$ so that for $x \sim \mathcal{D}$, $\sum_{e \in P_i} x_e v_e$ is γ-subgaussian for every $i \in V$ where $P_i \subseteq V$ are the edges on the path from the root to i.*

Proof. Let $N \in \mathbb{N}$ be a large enough parameter that we determine later. We replace each edge $e \in E$ by a path consisting of N many edges where the vectors assigned to edges are

$$v_e^{(\ell)} := (\mathbf{0}, \dots, \mathbf{0}, v_e, \mathbf{0}, \dots, \mathbf{0}) \in \mathbb{R}^{Nn} \quad \forall \ell = 1, \dots, N$$

Note that still $\|v_e^{(\ell)}\|_2 \leq 1$. We call the new tree $\mathcal{T}' = (V', E')$ and note that $|E'| = N \cdot |E|$. We define K as in (9.1) and by Proposition 158 we have

$$\gamma_{Nn}(K) \geq 1 - \frac{C_\delta^n}{N^{1+\delta}} \geq 1 - \frac{1}{2|E'|},$$

choosing $N := (2|E| \cdot C_\delta^n)^{1/\delta}$. We apply Theorem 157 to the tree $\mathcal{T}'$ in order to obtain signs $(x_e^\ell)_{e \in E, \ell \in [N]} \in \{-1, 1\}^{|E| \cdot N}$ so that

$$\sum_{\ell=1}^{N} \sum_{e \in P_i} x_e^{(\ell)} v_e^{(\ell)} \in \alpha K$$

for all $i \in V$ where $\alpha < 5$. If we draw an index $\ell \sim [N]$ uniformly and set $X := (x_1^{(\ell)}, \dots, x_{|E|}^{(\ell)})$, then by construction of K one has

$$\Big\| \sum_{e \in P_i} X_e v_e \Big\|_{\psi_2, \infty} \leq \alpha \cdot (2 + \delta) < 10 \ \ \forall i \in V$$

if we choose $\delta > 0$ small enough. □

Note that the produced distribution $\mathcal{D}$ is indeed the uniform distribution over a multiset of $(2|E| \cdot C_\delta^n)^{1/\delta}$ sign vectors. We remark that even if $\mathcal{T}$ is a path with the n vectors $e_1, \dots, e_n$, any $O(1)$-subgaussian distribution over sign vectors needs to have support at least $2^{\Omega(n)}$.

We should also point out that our proof strategy is related to an argument by Raghu Meka to use Banaszczyk's Theorem 11 to prove the existence of a good SDP solution. This was reported in the work of [26] in the context of an SDP solution with small discrepancy for all prefixes. After completing a preliminary draft, we learned the existence of an unpublished manuscript [125] proving Theorem 161 for the special case where $\mathcal{T}$ is a path.

9.6 Existence of an online algorithm

Finally, we use the tree subgaussianity from Theorem 161 to show the existence of an online algorithm which maintains subgaussian prefixes.

Theorem 162. *For every $T \in \mathbb{N}$, there exists a randomized online algorithm which, upon receiving a vector $v_i \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$ for each $i \in [T]$, outputs a random sign $x_i \in \{-1, 1\}$ so that the prefix sum $\sum_{j=1}^{i} x_j v_j$ is 10-subgaussian. The algorithm runs in time $\exp(T^{CnT})$ for some universal constant $C > 0$.*

Proof. We write the constant from Theorem 161 as $10 - \delta$ for some $\delta > 0$. Let W be an ε-net of S^{n-1} of size $(3/\varepsilon)^n$ where $\varepsilon := \frac{\delta}{10T}$. We consider a rooted tree $\mathcal{T} = (V, E)$ of depth T so that all non-leaf nodes have $|W|$ children, labeled with each of the vectors in W. By Theorem 161, there exists a distribution $\mathcal{D}$ over signs $x \in \{-1, 1\}^E$ that is $(10 - \delta)$-subgaussian.

We now describe the randomized online algorithm. First, we sample random signs $x \sim \mathcal{D}$ and keep track of a position $p \in V$ in the tree, initially set to be the root. Now for each incoming vector $v_i \in \mathbb{R}^n$, we output the sign corresponding to an edge (p, p') labelled with v_i' that satisfies $\|v_i - v_i'\|_2 \leq \varepsilon$ and set $p := p'$. It remains to argue that the resulting prefix sums are indeed 10-subgaussian. Indeed, for any i one has

$$\Big\| \sum_{j=1}^{i} x_j v_j \Big\|_{\psi_2,\infty} \leq \underbrace{\Big\| \sum_{j=1}^{i} x_j v_j' \Big\|_{\psi_2,\infty}}_{\leq 10-\delta} + \Big\| \sum_{j=1}^{i} x_j (v_j - v_j') \Big\|_{\psi_2,\infty}$$

$$\leq 10 - \delta + \sum_{j=1}^{i} \sup_{w \in S^{m-1}} \underbrace{|\langle v_j - v_j', w \rangle|}_{\leq \varepsilon} \cdot \underbrace{\|x_j\|_{\psi_2}}_{\leq 2}$$

$$\leq 10 - \delta + T \cdot 2\varepsilon < 10.$$

Next, we discuss the running time of this procedure. The total number of edges in $\mathcal{T}$ equals $|E| = |W| + \cdots + |W|^T \leq (2|W|)^T \leq O(T/\delta)^{nT}$. Reinspecting the proof of Theorem 161 we recall that the $(10 - \delta)$-subgaussian distribution $\mathcal{D}$ is constructed from a proper sign vector for the tree $\mathcal{T}' = (V', E')$ that has $|E'| = N \cdot |E|$ many edges where $N := (2|E| \cdot C_\delta^n)^{1/\delta}$. The number of candidate sign vectors to be tried out is then $2^{|E'|} \leq \exp((T/\delta)^{O(nT)})$, and we can verify the subgaussianity of each by computing the subgaussian norm of each root-vertex path sum over inner products with all vectors from W in time $\exp(T^{CnT})$; then Proposition 160 guarantees that if all such inner products are $(10-\delta)$-subgaussian, then the root-vertex path sums are also $\frac{10-\delta}{1-\varepsilon} \leq \frac{10-\delta}{1-\delta/10} = 10$-subgaussian. $\square$

9.7 *An online algorithm independent of the input length*

In the previous section we argued that for every number T of vectors, there is an online algorithm that balances T vectors. In this section, we prove that in fact, there has to be a *single* online algorithm that balances any sequence of vectors without knowing the number vectors beforehand. Note that just using Theorem 162 as a black box, the online algorithm to balance T vectors could be very different from the algorithm to balance $T' < T$ vectors. But of course one could have used the algorithm that worked for T vectors also to balance just T' vectors. By using a compactness argument we will argue that there is indeed a single algorithm.

We fix some $\varepsilon > 0$ and set $\varepsilon_i := \varepsilon 2^{-i}$ for all $i \geq 1$. Let $W_i \subseteq S^{n-1}$ be an ε_i-net. Let $\mathcal{T}_i = (V_i, E_i)$ be a tree of height i with a distinguished root r where for all $j \in \{1, \ldots, i\}$, each node at depth $j-1$ has $|W_j|$ many outgoing edges, one labelled with each vector from W_j. In other words, any root-leaf path in $\mathcal{T}_i$ corresponds to a sequence $(v_1, \ldots, v_i)$ with $v_j \in W_j$ for $j = 1, \ldots, i$. We write $\mathcal{T}^* = (V^*, E^*)$ as the infinite tree constructed in the same manner, i.e. for all $j \geq 1$, nodes at distance $j-1$ to the root have $|W_j|$ children. We say that a distribution $\mathcal{D}_i$ over signs $\Omega_i := \{-1, 1\}^{E_i}$ is *c-subgaussian for $\mathcal{T}_i$* if $\|\sum_P x_e v_e\|_{\psi_2} \leq c$ for $x \sim \mathcal{D}_i$ and every path P in $\mathcal{T}_i$ starting at the root. First we prove that the distributions for different height trees can be chosen in a consistent way:

Lemma 163. *There is a constant $\gamma < 10$ so that for every $\varepsilon > 0$ there is a family of distributions $\{\mathcal{D}_i^*\}_{i \geq 1}$ where each $\mathcal{D}_i^*$ is a distribution over Ω_i that is γ-subgaussian for $\mathcal{T}_i$ and $\mathcal{D}_i^* = \Pi_{\Omega_i}(\mathcal{D}_{i+1}^*)$*

for all $i \geq 1$.

Here $\Pi_{\Omega_i}(\mathcal{D}^*_{i+1})$ denotes the projection (or the marginals) of $\mathcal{D}^*_{i+1}$ on Ω_i.

Proof. Consider the sequence of γ-subgaussian distributions $\mathcal{D}_i$ for $\mathcal{T}_i$ (which depend on ε) given by Theorem 161 as points in the metric space $\Delta_i := \{x \in \mathbb{R}^{\Omega_i}_{\geq 0} : \|x\|_1 = 1\}$. Since Δ_i is closed and bounded, it is compact. For every $i' \leq i$, since $\mathcal{T}_{i'} \subseteq \mathcal{T}_i$, it follows that $\Pi_{\Omega_{i'}}(\mathcal{D}_i)$ is also γ-subgaussian.

We construct $\mathcal{D}^*_i$ inductively. First consider the infinite sequence of distributions $\Pi_{\Omega_1}(\mathcal{D}_i)$, all of which are γ-subgaussian for $\mathcal{T}_1$. Since Δ_i is compact, there exists a subsequence of indices $\{k_{j,1}\}_{j\geq 1}$ so that $\mathcal{D}^*_1 := \lim_{j\to\infty} \Pi_{\Omega_1}(\mathcal{D}_{k_{j,1}})$ exists. By continuity, $\mathcal{D}^*_1$ is also γ-subgaussian over $\mathcal{T}_1$.

Now assume that we have constructed distributions $\mathcal{D}^*_\ell$ for $1 \leq \ell \leq i$ that are γ-subgaussian for $\mathcal{T}_\ell$ with $\mathcal{D}^*_\ell = \Pi_{\Omega_\ell}(\mathcal{D}^*_{\ell+1})$ for $1 \leq \ell < i$, as well as an infinite sequence of indices $\{k_{j,i}\}_{j\geq 1}$ which again satisfy $\mathcal{D}^*_\ell = \lim_{j\to\infty} \Pi_{\Omega_\ell}(\mathcal{D}_{k_{j,i}})$ for $\ell \leq i$. We may drop a prefix if needed so that $k_{j,i} \geq i+1$ for all $j \geq 1$. By compactness of Δ_{i+1}, there exists a subsequence of indices $\{k_{j,i+1}\}_{j\geq 1}$ so that $\mathcal{D}^*_{i+1} := \lim_{j\to\infty} \Pi_{\Omega_{i+1}}(\mathcal{D}_{k_{j,i+1}})$ exists; the previous limits remain the same. Again by continuity of the subgaussian norm, $\mathcal{D}^*_{i+1}$ is γ-subgaussian over $\mathcal{T}_{i+1}$, and also $\Pi_{\Omega_i}(\mathcal{D}^*_{i+1}) = \Pi_{\Omega_i}(\lim_{j\to\infty} \Pi_{\Omega_{i+1}}(\mathcal{D}_{k_{j,i+1}})) = \lim_{j\to\infty} \Pi_{\Omega_i}(\mathcal{D}_{k_{j,i+1}}) = \mathcal{D}^*_i$. $\square$

Theorem 164. *There exists an online algorithm which, upon receiving a vector $v_i \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$, outputs a random sign $x_i \in \{-1, 1\}$ so that the prefix sum $\sum_{j=1}^i x_j v_j$ is 10-subgaussian.*

Proof. Fix a small enough $\delta > 0$ so that Lemma 163 works with constant $10 - \delta$. Let $\mathcal{D}^*_i$ be $(10 - \delta)$-subgaussian distributions provided by Lemma 163 where we choose $\varepsilon := \frac{\delta}{4}$. We keep track of a position $p \in V$ in the tree, initially set to be the root. For each incoming vector $v_i \in \mathbb{R}^n$, we sample a random sign x_i from $\mathcal{D}^*_i$ *conditioned* on the previous signs $x_1, \ldots, x_{i-1}$ on edges already visited, corresponding to an edge (p, p') labeled with a vector v'_i that satisfies $\|v_i - v'_i\|_2 \leq \varepsilon_i$, and set $p := p'$. It remains to argue that the resulting prefix sums are 10-subgaussian. Indeed, since $\mathcal{D}^*_j = \Pi_{\Omega_j}(\mathcal{D}^*_i)$ for all $j < i$, the distribution of the signs $(x_1, \ldots, x_i)$ is equivalent to drawing all of them from $\mathcal{D}_i^*$. Then, as in Theorem 162,

$$\begin{aligned}\Big\| \sum_{j=1}^{i} x_j v_j \Big\|_{\psi_2,\infty} &\leq \underbrace{\Big\| \sum_{j=1}^{i} x_j v'_j \Big\|_{\psi_2,\infty}}_{\leq 10-\delta} + \Big\| \sum_{j=1}^{i} x_j (v_j - v'_j) \Big\|_{\psi_2,\infty} \\ &\leq 10 - \delta + \sum_{j=1}^{i} \sup_{w \in S^{m-1}} \underbrace{|\langle v_j - v'_j, w\rangle|}_{\leq \varepsilon_j} \cdot \underbrace{\|x_j\|_{\psi_2}}_{\leq 2} \\ &\leq 10 - \delta + 2\varepsilon < 10. \end{aligned}$$

□

9.8 *Applications*

In this section, we show a few direct consequences of our main Theorem 145.

Proof of Theorem 146. Both items will follow from Lemma 156. For the first item, note that by Lemmas 26 and 27 in [51], it follows that any symmetric convex body K with $\gamma_n(K) \geq 1/2$ satisfies $E_{g\sim N(\mathbf{0},I_n)}[\|g\|_K] \lesssim 1$ and $\mathrm{inradius}(K) \gtrsim 1$, so it suffices to apply Lemma 156 with $\delta := \frac{1}{2}$.

For the second item, any symmetric convex body K with $\gamma_n(K) \geq 1 - \frac{1}{2T}$ once again satisfies $\mathbb{E}_{g\sim N(\mathbf{0},I_n)}[\|g\|_K] \lesssim 1$, and moreover K must contain a $\Omega(\sqrt{\log T})$-radius ball, for otherwise it would be contained in a strip of Gaussian measure less than $1 - \frac{1}{2T}$. Therefore $\max_{x\in K^\circ} \|x\|_2 \lesssim 1/\sqrt{\log T}$ and we may apply Lemma 156 with $\delta := \frac{1}{2T}$ together with the union bound over the T prefixes. □

Proof of Corollary 147. Let $d \leq \min(n, T)$ denote the dimension of the linear span U of $v_1, \ldots, v_T$. By Corollary 19 in [30],

$$\mathop{\mathbb{E}}_{g\sim N(\mathbf{0},I_n)}[\|g\|_{U\cap B_p^n}] \leq \mathop{\mathbb{E}}_{g\sim N(\mathbf{0},I_d)}[\|g\|_{B_p^d}] \overset{\text{Jensen}}{\leq} \mathop{\mathbb{E}}_{g\sim N(\mathbf{0},I_d)}[\|g\|_{B_p^d}^p]^{1/p} \leq \sqrt{p} \cdot d^{1/p}.$$

Here we also used that $\mathbb{E}_{g\sim N(0,1)}[g^p] \leq p^{p/2}$. By Lemma 156,

$$\begin{aligned}\|X\|_p = \|X\|_{U\cap B_p^n} &\lesssim \mathop{\mathbb{E}}_{g\sim N(\mathbf{0},I_n)}[\|g\|_{U\cap B_p^n}] + \sqrt{\log(1/\delta)} \cdot \underbrace{1/\mathrm{inradius}(B_p^n)}_{=1} \\ &\leq \sqrt{p}\min(n,T)^{1/p} + \sqrt{\log(1/\delta)}\end{aligned}$$

with probability at least $1-\delta$, as claimed. The second item follows from a union bound over T such events and that $\sqrt{\log(T/\delta)} \leq \sqrt{\log T} + \sqrt{\log(1/\delta)}$.

For the ℓ_∞ bounds, we instead apply Theorem 9 in [30] which together with Lemma 26 in [51] gives $\mathbb{E}_{g\sim N(\mathbf{0},I_n)}[\|g\|_{U\cap B_\infty^n}] \lesssim \sqrt{\log\min(n,T)}$, so that Lemma 156 analogously yields both claims. $\square$

Proof of Corollary 149. It suffices to construct a vector $v_e := e_u - e_v \in \mathbb{R}^{|V|}$ for every edge $e = \{u, v\}$; then signs correspond to an orientation. Since such vectors have constant ℓ_2 norm, the claim follows directly from Corollary 147(d). $\square$

9.9 Online discrepancy lower bound

We show that the bound in Corollary 147(d) is tight up to a constant:

Proof of Theorem 148. We may also assume that $n = 2$ as otherwise we may pad the construction with zeros. By Yao's minimax principle [170], it suffices to show a strategy for the adversary against deterministic online algorithms.

Split the time horizon into T/k blocks of length $k := \lceil \frac{1}{2}\log T\rceil$. In each block, the adversary samples signs $y \sim \{-1,1\}^k$ uniformly at random, and outputs unit vectors $v_1, \ldots, v_k \in \mathbb{R}^2$ so that for each $i \in [k]$, v_i is orthogonal to $\sum_{j<i} y_j v_j$. The online algorithm then outputs deterministic signs $x \in \{-1,1\}^k$ where x_i only depends on y_j, v_j for $j < i$. Since the sign y_i is independent of the sign x_i given by the online algorithm, the sign vectors x and y will be equal with probability 2^{-k}. In this case, $\|\sum_{i=1}^k x_i v_i\|_2 = \sqrt{k}$. The probability that this happens in *at least one* block is $1-(1-2^{-k})^{T/k} = 1-2^{-T^{\Omega(1)}}$. Then one of the two prefixes induced by that block will have ℓ_2 norm at least $\frac{1}{2}\sqrt{k}$ and ℓ_∞ norm at least $\frac{1}{\sqrt{2}}\cdot\frac{1}{2}\sqrt{k} \gtrsim \sqrt{\log T}$. $\square$

9.10 Open problems

We mention two other settings for which the optimal discrepancy bound against oblivious adversaries remains open:

Conjecture 17. *Does there exist an online algorithm that for any sequence of vectors* $v_1, \ldots, v_n \in \mathbb{R}^n$ *with* $\|v_i\|_\infty \leq 1$*, arriving one at a time, decides random signs* $x_1, \ldots, x_n \in \{-1, 1\}$ *so that* $\|\sum_{i=1}^n x_i v_i\|_\infty \leq O(\sqrt{n})$ *with high probability?*

Corollary 147(c) gives a bound of $O(\sqrt{n \log n})$. Bansal and Spencer have settled the case where the vectors are chosen uniformly at random from $\{-1, 1\}^n$ [27]. Note that there is a $\Omega(\sqrt{n})$ lower bound even in the offline setting and a $\Omega(\sqrt{\log T})$ lower bound from Theorem 148 (thus we restrict to $T = n$).

Conjecture 18. *Does there exist an online algorithm that for any sequence of vectors* $v_1, \ldots, v_T \in \mathbb{R}^n$*, each with two nonzero coordinates (one equal to 1 and the other -1) and arriving one at a time, decides random signs* $x_1, \ldots, x_T \in \{-1, 1\}$ *so that* $\|\sum_{i=1}^t x_i v_i\|_\infty \leq O(\sqrt[3]{\log T})$ *for all* $t \in [T]$ *with high probability?*

Corollary 149 gives an upper bound of $O(\sqrt{\log T})$ and there is a $\Omega(\sqrt[3]{\log T})$ lower bound [2, 62].

Finally, we ask for a polynomial time algorithm for Theorem 145.

Conjecture 19. *Does there exist a polynomial time online algorithm that against any oblivious adversary, for any sequence of vectors* $v_1, \ldots, v_T \in \mathbb{R}^n$ *with* $\|v_i\|_2 \leq 1$*, decides random signs* $x_1, \ldots, x_T \in \{-1, 1\}$ *so that for every* $t \in [T]$*, the prefix sum* $\sum_{i=1}^t x_i v_i$ *is* $O(1)$*-subgaussian?*

Chapter 10

THE SUBSPACE FLATNESS CONJECTURE AND FASTER INTEGER PROGRAMMING

This chapter is based on joint work with Thomas Rothvoss [141]; the improved bound improved bound in Theorem 171 is due to Daniel Dadush.

10.1 Introduction

Lattices are fundamental objects studied in various areas of mathematics and computer science. Here, a *lattice* Λ is a discrete subgroup of $\mathbb{R}^n$. If $B \in \mathbb{R}^{n\times k}$ is a matrix with linearly independent columns $b_1, \dots, b_k$, then we may write a lattice in the form $\Lambda(B) := \{\sum_{i=1}^k y_i b_i : y_i \in \mathbb{Z}\}$. In mathematics, lattices are the central object of study in the geometry of numbers with many applications for example to number theory, see e.g. [83]. On the computer science side, lattices found applications for example in lattice-based cryptography [136] and cryptanalysis [127]. One of the most important algorithms at least in this area is the *LLL-algorithm* by Lenstra, Lenstra and Lovász [98] which finds an approximately orthogonal basis for a given lattice in polynomial time. One of the consequences of the LLL-reduction is a polynomial time $2^{n/2}$-approximation algorithm for the problem of finding a (nonzero) *shortest vector* in a lattice. We should also mention that the problem of finding a shortest vector in any norm can be solved in time $2^{O(n)}$ using a variation of the sieving algorithm [3] while in the Euclidean norm, even the closest vector to any given target vector can be found in time $2^{O(n)}$ [117]. A more general problem with tremendous applications in combinatorial optimization and operations research is the one of finding an integer point in an arbitrary convex body or polytope. Lenstra [99] used the then-recent lattice basis reduction algorithm to solve any n-variable integer program in time

$2^{O(n^2)}$. This was later improved by Kannan [82] to $n^{O(n)}$ and then by Dadush [58] and by Dadush, Eisenbrand and Rothvoss [50] to $2^{O(n)}n^n$.

A parameter appearing in the geometry of numbers is the *covering radius*

$$\mu(\Lambda, K) := \min\left\{r \geq 0 \mid \Lambda + rK = \text{span}(\Lambda)\right\}$$

of a lattice $\Lambda \subseteq \mathbb{R}^n$ with respect to a compact convex set $K \subseteq \mathbb{R}^n$ with $\text{span}(\Lambda) = \text{affine.hull}(K)$. This quantity seems to be substantially harder computationally, in the sense that the question whether $\mu(\Lambda, K)$ is at least/at most a given threshold seems to be neither in **NP** nor in **coNP**. In terms of approximating $\mu(\Lambda, K)$, one can quickly observe that one has the lower bound of $\mu(\Lambda, K) \geq (\frac{\det(\Lambda)}{\text{Vol}_n(K)})^{1/n}$, simply because for $r < (\frac{\det(\Lambda)}{\text{Vol}_n(K)})^{1/n}$, the average density of the translates $\Lambda + rK$ is less than 1. However, this lower bound may be arbitrarily far off the real covering radius, for example if $\Lambda = \mathbb{Z}^2$ and $K = [-\frac{1}{M}, \frac{1}{M}] \times [-M, M]$ with $M \to \infty$. On the other hand, for any subspace $W \subseteq \mathbb{R}^n$ one trivially has $\mu(\Lambda, K) \geq \mu(\Pi_W(\Lambda), \Pi_W(K))$, where Π_W is the orthogonal projection into W. Hence, following Kannan and Lovász [83], one might instead consider the best volume based lower bound for any projection, i.e.

$$\mu_{KL}(\Lambda, K) := \max_{\substack{W \subseteq \text{span}(\Lambda) \text{ subspace} \\ d := \dim(W)}} \Big(\frac{\det(\Pi_W(\Lambda))}{\text{Vol}_d(\Pi_W(K))}\Big)^{1/d}$$

Kannan and Lovász [83] indeed provide an upper bound of

$$\mu_{KL}(\Lambda, K) \leq \mu(\Lambda, K) \leq n \cdot \mu_{KL}(\Lambda, K)$$

On the other hand, they also construct a simplex $K \subseteq \mathbb{R}^n$ for which $\mu(\mathbb{Z}^n, K) \geq \Omega(\log(2n)) \cdot \mu_{KL}(\mathbb{Z}^n, K)$ holds. Dadush [58] states the following conjecture, attributing it to Kannan and Lovász [83]:

Conjecture 20 (Subspace Flatness Conjecture)**.** *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$ one has*

$$\mu_{KL}(\Lambda, K) \leq \mu(\Lambda, K) \leq O(\log(2n)) \cdot \mu_{KL}(\Lambda, K)$$

Dadush also realized the tremendous implications of this conjecture to optimization and showed that it would imply a $O(\log(2n))^n$-time algorithm to solve n-variable integer programs, assuming that the subspace W attaining $\mu_{KL}(\Lambda, K)$ could also be found in the same time. Later, Dadush and Regev [55] conjectured a *Reverse Minkowski-type Inequality*, which intuitively says that any lattice without dense sublattices should contain only few short vectors. Among other applications, they proved that this conjecture would imply Conjecture 20 (with some logarithmic loss) at least for the case that K is an ellipsoid. The conjecture of [55] was then resolved by Regev and Stephens-Davidowitz [137] with a rather ingenious proof. More precisely, they prove the following:

Theorem 165 (Reverse Minkowski Theorem [137])**.** *Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice that satisfies $\det(\Lambda') \geq 1$ for all sublattices $\Lambda' \subseteq \Lambda$. Then for a large enough constant $C > 0$ and $s = C\log(2n)$ one has $\rho_{1/s}(\Lambda) \leq \frac{3}{2}$.*

Here, one has $\rho_t(x) := \exp(-\pi\|x/t\|_2^2)$ where $t > 0$ and for a discrete set $S \subseteq \mathbb{R}^n$ we abbreviate $\rho_t(S) := \sum_{x \in S} \rho_t(x)$. To understand the power of this result compared to classical arguments, note that from $\det(\Lambda') \geq 1$ for all $\Lambda' \subseteq \Lambda$ one can derive that each vector $x \in \Lambda \setminus \{\mathbf{0}\}$ has length $\|x\|_2 \geq 1$ and so by a standard packing argument we know that for any $r \geq 1$ one has $|\Lambda \cap rB_2^n| \leq (3r)^n$, which is exponential in n. On the other hand, again under the assumption that $\det(\Lambda') \geq 1$ for all $\Lambda' \subseteq \Lambda$, the Reverse Minkowski Theorem implies that $|\Lambda \cap rB_2^n| \leq \exp(\Theta(\log^2(2n)) \cdot r^2)$ which is quasi-polynomial in n. Also, [137] tighten the reduction to the Subspace Flatness Conjecture and show that it holds for any ellipsoid with a factor of $O(\log^{3/2}(2n))$. While for any convex body K, there is an ellipsoid $\mathcal{E}$ and a center c so that $c + \mathcal{E} \subseteq K \subseteq c + n\mathcal{E}$ [80], this factor of n is the best possible, and hence there does not seem to be a blackbox reduction from the general case of Conjecture 20 to the one of ellipsoids.

10.1.1 Our contribution

Our main result is as follows:

Theorem 166. *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$ one has*

$$\mu_{KL}(\Lambda, K) \leq \mu(\Lambda, K) \leq O(\log^3(2n)) \cdot \mu_{KL}(\Lambda, K).$$

We will break the proof into two parts that can be found in Section 10.4. Our result is constructive in the following sense:

Theorem 167. *Given a full rank lattice $\Lambda := \Lambda(B)$ and a convex body $K \subseteq \mathbb{R}^n$ with $c + r_0 B_2^n \subseteq K \subseteq r_1 B_2^n$, there is a randomized algorithm to find a subspace $W \subseteq \mathbb{R}^n$ with $d := \dim(W)$ so that*

$$\mu(\Lambda, K) \leq O(\log^4(2n)) \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d}.$$

The running time of that algorithm is $2^{O(n)}$ times a polynomial in $\log(\frac{1}{r_0})$, $\log(r_1)$ and in the encoding length of B.

Here, a separation oracle suffices for K. See Section 10.5 for a proof. Following the framework layed out by Dadush [58], this implies a faster algorithm to find a lattice point in a convex body:

Theorem 168. *Given a convex body $K \subseteq rB_2^n$ represented by a separation oracle and a lattice $\Lambda = \Lambda(B)$, there is a randomized algorithm that with high probability finds a point in $K \cap \Lambda$ or correctly decides that there is none. The running time is $(\log(2n))^{O(n)}$ times a polynomial in $\log(r)$ and the encoding length of B.*

The proof can be found in Section 10.6. Applying Theorem 168 to integer programming we obtain the following:

Theorem 169. *Given $A \in \mathbb{Q}^{m \times n}$, $b \in \mathbb{Q}^m$ and $c \in \mathbb{Q}^n$, the integer linear program $\max\{c^T x \mid Ax \leq b, x \in \mathbb{Z}^n\}$ can be solved in time $(\log(2n))^{O(n)}$ times a polynomial in the encoding length of A, b and c.*

An immediate consequence of our main result (Theorem 166) is that K can be replaced by a larger *symmetric* body without decreasing the covering radius significantly:

Theorem 170. *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$ one has*

$$\mu(\Lambda, K - K) \leq \mu(\Lambda, K) \leq O(\log^3(2n)) \cdot \mu(\Lambda, K - K).$$

Another consequence is that the *flatness constant* in dimension n is bounded by $O(n \log^3(2n))$, which is an improvement from the previously known bound of $O(n^{4/3} \log^{O(1)}(2n))$ obtained by combining the result of Rudelson [148] with [15].

Theorem 171. *For any convex body $K \subseteq \mathbb{R}^n$ and any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ one has*

$$\mu(\Lambda, K) \cdot \lambda_1(\Lambda^*, (K - K)^\circ) \leq O(n \log^3(2n)).$$

It is well known that Theorem 171 can also be rephrased in the following convenient form:

Corollary 172. *Let $K \subseteq \mathbb{R}^n$ by a convex body with $K \cap \mathbb{Z}^n = \emptyset$. Then there is a vector $c \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$ so that at most $O(n \log^3(2n))$ many hyperplanes of the form $\langle c, x\rangle = \delta$ with $\delta \in \mathbb{Z}$ intersect K.*

We will prove Theorem 170, Theorem 171 and Corollary 172 in Section 10.7.

10.2 Preliminaries

In this section, we introduce the tools that we rely on later. We write $A \lesssim B$ if there is a universal constant $C > 0$ so that $A \leq C \cdot B$ holds. We write $A \asymp B$ if both $A \lesssim B$ and $B \lesssim A$ hold.

10.2.1 Lattices

For a lattice $\Lambda = \Lambda(B)$ given by a matrix $B \in \mathbb{R}^{n \times k}$ with linearly independent columns, we define the *rank* as $\mathrm{rank}(\Lambda) := k = \dim(\mathrm{span}(\Lambda))$ and the *determinant* as $\det(\Lambda) = \sqrt{\det_k(B^T B)}$. A lattice $\Lambda \subseteq \mathbb{R}^n$ with $\mathrm{rank}(\Lambda) = n$ has *full rank*. For a lattice $\Lambda \subseteq \mathbb{R}^n$, we define the *dual lattice* as $\Lambda^* := \{x \in \mathrm{span}(\Lambda) \mid \langle x, y\rangle \in \mathbb{Z}\ \forall y \in \Lambda\}$. Recall that $\det(\Lambda) \cdot \det(\Lambda^*) = 1$. A consequence of the *Poisson Summation Formula* is as follows:

Lemma 173. *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$, vector $u \in \mathbb{R}^n$ and any $s > 0$ one has*

$$|\rho_s(\Lambda + u) - s^n \det(\Lambda^*)| \leq s^n \det(\Lambda^*) \cdot \rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}).$$

A set $K \subseteq \mathbb{R}^n$ is called a *convex body* if it is convex, compact (i.e. bounded and closed) and has a non-empty interior int(K). A set Q is called *symmetric* if $-Q = Q$. For a symmetric convex set Q, the norm $\|x\|_Q$ is defined as the least scaling $r \geq 0$ so that $x \in rQ$. For a lattice Λ and a symmetric convex body Q we denote the length of the shortest vector by

$$\lambda_1(\Lambda, Q) := \min_{x \in \Lambda \setminus \{\mathbf{0}\}} \|x\|_Q$$

Later we will also need a classical bound on short vectors in a lattice:

Theorem 174 (Minkowski's First Theorem)**.** *Let $\Lambda \subseteq \mathbb{R}^n$ be a full rank lattice and $Q \subseteq \mathbb{R}^n$ be a symmetric convex body. Then $\lambda_1(\Lambda, Q) \leq 2\Big(\frac{\det(\Lambda)}{\mathrm{Vol}_n(Q)}\Big)^{1/n}$.*

We recommend the excellent notes of Regev [135] for background.

10.2.2 Stable lattices and the canonical filtration

A subspace $W \subseteq \mathbb{R}^n$ is a *lattice subspace* of a lattice $\Lambda \subseteq \mathbb{R}^n$ if span$(W \cap \Lambda) = W$. Similarly, a sublattice $\Lambda' \subseteq \Lambda$ is called *primitive* if there is a subspace W with $\Lambda \cap W = \Lambda'$. For a lattice Λ and a primitive sublattice $\Lambda' \subseteq \Lambda$, we define the *quotient lattice* as $\Lambda/\Lambda' := \Pi_{\text{span}(\Lambda')^\perp}(\Lambda)$. In many ways one can imagine that the quotient operation factors Λ into two lattices Λ' and Λ/Λ'. In particular Λ' and Λ/Λ' are orthogonal and $\det(\Lambda) = \det(\Lambda') \cdot \det(\Lambda/\Lambda')$.

A lattice $\Lambda \subseteq \mathbb{R}^n$ is called *stable* if $\det(\Lambda) = 1$ and $\det(\Lambda') \geq 1$ for all sublattices $\Lambda' \subseteq \Lambda$. That means a stable lattice does not contain any sublattice that is denser than the lattice itself. One can easily verify that for example $\mathbb{Z}^n$ is stable. We denote nd$(\Lambda) := \det(\Lambda)^{1/\text{rank}(\Lambda)}$ as the *normalized determinant*. One can prove that the extreme points of the 2-dimensional convex hull of the points $\big\{(\text{rank}(\Lambda'), \ln(\det(\Lambda'))) \mid \text{sublattice } \Lambda' \subseteq \Lambda\big\}$ correspond to a unique chain of nested sublattices $\{\mathbf{0}\} = \Lambda_0 \subset \Lambda_1 \subset \ldots \subset \Lambda_k = \Lambda$. That chain is called the *canonical filtration*. It is useful to observe that each Λ_i in this sequence is the unique densest sublattice of Λ with given dimension $\mathrm{rank}(\Lambda_i)$. Moreover, the quotient lattices Λ_i/Λ_{i-1} are all scalars of a stable lattice and one can prove that $\mathrm{nd}(\Lambda_i/\Lambda_{i-1})$ are strictly increasing in i. We refer to the thesis of [160] for details.

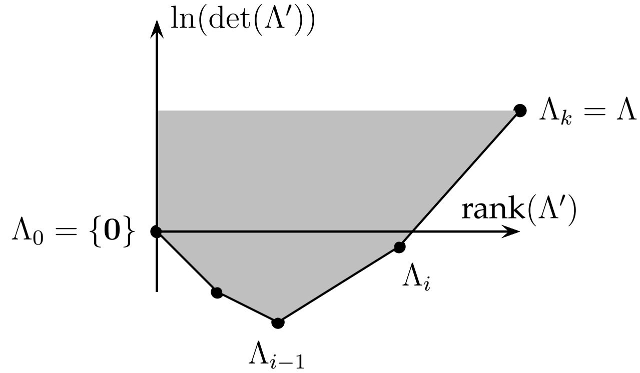

It will be useful to replace the canonical filtration by an *approximate* filtration where the normalized determinants grow exponentially. We make the following definition:

Definition 175. *We call a lattice $\Lambda \subseteq \mathbb{R}^n$ t-*stable *with $t \geq 1$ if the following holds:*

(*I*) *For any sublattice $\tilde{\Lambda} \subseteq \Lambda$ one has $\mathrm{nd}(\tilde{\Lambda}) \geq t^{-1}$.*

(*II*) *For any sublattice $\tilde{\Lambda} \subseteq \Lambda^*$ one has $\mathrm{nd}(\tilde{\Lambda}) \geq t^{-1}$.*

Note that a lattice is 1-stable if and only if it is stable. We can similarly define t-stable filtrations:

Definition 176. *Given a lattice $\Lambda \subseteq \mathbb{R}^n$, we call a sequence $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$ a t-*stable filtration *of Λ if the following holds:*

(*a*) *The normalized determinants $r_i := \mathrm{nd}(\Lambda_i/\Lambda_{i-1})$ satisfy $r_1 < \ldots < r_k$.*

(*b*) *The lattices $\frac{1}{r_i}(\Lambda_i/\Lambda_{i-1})$ are t-stable for all $i = 1, \ldots, k$.*

We call a t-stable filtration well-separated *if additionally the following holds:*

(*c*) *One has $r_i \leq \frac{1}{2} r_{i+2}$ for all $i = 1, \ldots, k-2$.*

For example, the canonical filtration is 1-stable. It turns out we can make any t-stable filtration well-separated:

Theorem 177. *Given a lattice $\Lambda \subseteq \mathbb{R}^n$ and a t-stable filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$, in polynomial time we can compute a $2t$-stable well-separated filtration $\{\mathbf{0}\} = \tilde{\Lambda}_0 \subseteq \ldots \subseteq \tilde{\Lambda}_{\tilde{k}} = \Lambda$.*

We defer the proof to Appendix 10.8. Using the canonical filtration as input to Theorem 177 yields:

Corollary 178. *For any lattice $\Lambda \subseteq \mathbb{R}^n$, there exists a 2-stable well-separated filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$.*

We collect a few more properties of t-stable lattices:

Lemma 179. *There is a universal constant $C > 0$ so that the following holds: Let Λ be a t-stable lattice for $t \geq 1$. Then for $s = C\log(2n)$ one has*

(a) *Λ^* is t-stable.*

(b) $\rho_{1/(st)}(\Lambda) \leq \frac{3}{2}$.

(c) *For any $u \in \mathbb{R}^n$ one has $\frac{\rho_{st}(\Lambda+u)}{\rho_{st}(\Lambda)} \geq \frac{1}{3}$.*

Proof. (a) is immediate from the definition of t-stability. Next, let $s = C\log(2n)$ be the parameter from Theorem 165. For (b), we can see that for any $\Lambda' \subseteq t\Lambda$ one has $\det(\Lambda') \geq 1$ and so the Reverse Minkowski Theorem (Theorem 165) applies to the lattice $t\Lambda$. Then $\rho_{1/(st)}(\Lambda) = \rho_{1/s}(t\Lambda) \leq \frac{3}{2}$ which gives (b). For (c), applying Lemma 173 twice gives

$$\frac{\rho_{st}(\Lambda+u)}{\rho_{st}(\Lambda)} \geq \frac{(st)^n \det(\Lambda^*) \cdot (1 - \rho_{1/(st)}(\Lambda^* \setminus \{\mathbf{0}\}))}{(st)^n \det(\Lambda^*) \cdot (1 + \rho_{1/(st)}(\Lambda^* \setminus \{\mathbf{0}\}))} \overset{(a)+(b)}{\geq} \frac{1 - \frac{1}{2}}{1 + \frac{1}{2}} = \frac{1}{3}. \qquad \square$$

10.2.3 The ℓ-value and volume estimates

We review a few results from convex geometry that can all be found in the textbook by Artstein-Avidan, Giannopoulos and Milman [7]. We denote $B_2^n := \{x \in \mathbb{R}^n \mid \|x\|_2 \leq 1\}$

and $S^{n-1} := \{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$ as the Euclidean ball and sphere, resp. Let $\nu_n := \mathrm{Vol}_n(B_2^n)$. The *relative interior* of K is $\mathrm{rel.int}(K) := \{x \in K \mid \exists \varepsilon > 0 : (x + \varepsilon \cdot B_2^n) \cap \mathrm{affine.hull}(K) \subseteq K\}$.

We define the *mean width* of a convex body K as $w(K) := \mathbb{E}_{\theta \sim S^{n-1}}[\max\{\langle \theta, x - y\rangle : x, y \in K\}]$. For a compact convex $K \subseteq \mathbb{R}^n$ with $\mathbf{0} \in \mathrm{rel.int}(K)$ we denote its *polar* by $K^\circ := \{y \in \mathrm{span}(K) : \langle x, y\rangle \leq 1 \ \forall x \in K\}$. Recall the following basic facts.

Lemma 180 (Properties of polarity)**.** *For two convex bodies $K, Q \subseteq \mathbb{R}^n$ with $\mathbf{0} \in int(K)$ and $\mathbf{0} \in int(Q)$ the following holds:*

(*a*) *One has $(K^\circ)^\circ = K$.*

(*b*) *For any subspace $F \subseteq \mathbb{R}^n$ one has $\Pi_F(K)^\circ = K^\circ \cap F$.*

(*c*) *One has $(K \cap Q)^\circ = conv(K^\circ \cup Q^\circ)$.*

(*d*) *One has $(-K)^\circ = -K^\circ$.*

We write $N(\mathbf{0}, I_n)$ as the standard Gaussian distribution on $\mathbb{R}^n$. The *ℓ-value* of a symmetric convex $Q \subseteq \mathbb{R}^n$ is defined as

$$\ell_Q = \mathop{\mathbb{E}}_{x \sim N(\mathbf{0}, I_n)}[\|x\|_Q^2]^{1/2}$$

One may think of ℓ_Q as the "average thinness" of Q. It turns out that the ℓ-value is also related to the mean width. To see this, note that $\|\cdot\|_{Q^\circ}$ is the *dual norm* to $\|\cdot\|_Q$, i.e. for all $x \in \mathbb{R}^n$ one has $\|x\|_{Q^\circ} = \max\{\langle x, y\rangle : y \in Q\}$. Then

$$\ell_{Q^\circ} = \mathop{\mathbb{E}}_{x \sim N(\mathbf{0}, I_n)}[\|x\|_{Q^\circ}^2]^{1/2} = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim N(\mathbf{0}, I_n)}\Big[\max\{\langle x, y\rangle^2 : y \in Q\}\Big]^{1/2} \tag{10.1}$$

We can see that the right hand side of (10.1) almost matches the definition of $w(Q)$. In fact, one can prove:

Lemma 181. *For any symmetric convex body $Q \subseteq \mathbb{R}^n$ one has $\ell_{Q^\circ} \asymp \sqrt{n} \cdot w(Q)$.*

For a positive semidefinite matrix Σ we write $N(\mathbf{0},\Sigma)$ as the Gaussian with mean $\mathbf{0}$ and covariance matrix Σ and for a subspace $U \subseteq \mathbb{R}^n$ we write I_U as the identity matrix on that subspace. Occasionally we will need to refer to the ℓ-value of a compact symmetric convex set Q that is not necessarily full-dimensional. In that case we extend the definition to $\ell_Q = \mathbb{E}_{x\sim N(\mathbf{0},I_{\mathrm{span}(Q)})}[\|x\|_Q^2]^{1/2}$.

We say that a symmetric convex body Q is in *ℓ-position* if $\ell_Q \cdot \ell_{Q^\circ} \lesssim n\log(2n)$. One of the most powerful tools in convex geometry is that every symmetric convex body can indeed be brought into ℓ-position:

Theorem 182 (Figiel, Tomczak-Jaegerman, Pisier)**.** *For any symmetric convex body $Q \subseteq \mathbb{R}^n$, there is an invertible linear map $T : \mathbb{R}^n \to \mathbb{R}^n$ so that $\ell_{T(Q)} \cdot \ell_{(T(Q))^\circ} \lesssim n\log(2n)$.*

By Lemma 181, the conclusion of Theorem 182 is equivalent to $w(T(Q)) \cdot w(T(Q)^\circ) \lesssim \log(2n)$. Moreover one can prove that for any symmetric convex body Q one has $w(Q) \cdot w(Q^\circ) \gtrsim w(B_2^n)^2 \gtrsim 1$. Then one can interpret Theorem 182 that every symmetric convex body can be linearly transformed so that in terms of mean width and average thinness it is within a $O(\log(2n))$-factor of the Euclidean ball. For the sake of comparison, we note that the bound that could be obtained via the more classical John's Theorem [80] would be of the order of $\sqrt{n}$. We would like to point out that Theorem 182 is only known for symmetric convex bodies, and it is open to what extent it generalizes to the non-symmetric case.

We state two estimates concerning monotonicity of the ℓ-value that will be crucial for our later arguments:

Lemma 183. *Let $Q \subseteq \mathbb{R}^n$ be a symmetric convex body. Then for any subspace $U \subseteq \mathbb{R}^n$, one has $\ell_{Q\cap U} \leq \ell_Q$.*

Proof. Indeed, one has

$$\ell_Q^2 = \mathbb{E}_{z\sim N(\mathbf{0},I_U)}\left[\mathbb{E}_{y\sim N(\mathbf{0},I_{U^\perp})}[\|z+y\|_Q^2]\right] \geq \mathbb{E}_{z\sim N(\mathbf{0},I_U)}\Big[\Big\|z + \underbrace{\mathbb{E}_{y\sim N(\mathbf{0},I_{U^\perp})}[y]}_{=\mathbf{0}}\Big\|_Q^2\Big] = \ell_{Q\cap U}^2,$$

where the inequality follows from Jensen's inequality and the convexity of $y \mapsto \|z+y\|_Q^2$.

□

Lemma 184. *Let $Q \subseteq \mathbb{R}^n$ be a symmetric convex body. For any subspaces $V \subset W \subseteq \mathbb{R}^n$, one has $\ell_{\Pi_{V^\perp}(Q \cap W)} \leq \ell_Q$.*

Proof. We have $\ell_{\Pi_{V^\perp}(Q\cap W)} \leq \ell_{Q\cap W\cap V^\perp} \leq \ell_Q$ using that $\Pi_{V^\perp}(Q \cap W) \supseteq Q \cap W \cap V^\perp$ and using Lemma 183. □

A slight variant of Urysohn's Inequality (Theorem 18) will be handy for us:

Corollary 185 (Urysohn Inequality II)**.** *For any symmetric convex body $Q \subseteq \mathbb{R}^n$ one has $\mathrm{Vol}_n(Q)^{1/n} \lesssim \frac{\ell_{Q^\circ}}{n}$.*

Proof. Applying Urysohn's Inequality I we obtain

$$\mathrm{Vol}_n(Q)^{1/n} \overset{\text{Thm 18}}{\lesssim} w(Q) \cdot \underbrace{\mathrm{Vol}_n(B_2^n)^{1/n}}_{\lesssim 1/\sqrt{n}} \overset{\text{Lem 181}}{\lesssim} \frac{\ell_{Q^\circ}}{n}$$

Here we use in particular that $\mathrm{Vol}_n(B_2^n) \leq (\frac{2e}{\sqrt{n}})^n$. □

The following can be found e.g. in [7], Chapter 8:

Theorem 186 (Blaschke-Santaló-Bourgain-Milman)**.** *For any symmetric convex body $K \subseteq \mathbb{R}^n$ one has*

$$C_1^n \nu_n^2 \leq \mathrm{Vol}_n(K) \cdot \mathrm{Vol}_n(K^\circ) \leq C_2^n \nu_n^2$$

where $C_1, C_2 > 0$ are constants.

Let $b(K) := \frac{1}{\mathrm{Vol}_n(K)} \int_K x \, dx$ denote the *barycenter* or *centroid* of a convex body K. We will run into the issue that we need to control the volume of a non-symmetric convex body K, but Theorem 182 only holds for symmetric ones. A popular strategy in convex geometry is to translate K so that $b(K) = \mathbf{0}$ and then consider the *inner symmetrizer* $K \cap -K$ which by construction is a symmetric convex body contained in K which captures much of the geometry of K. For example a classical result by Milman and Pajor says that $\mathrm{Vol}_n(K \cap -K) \geq 2^{-n}\mathrm{Vol}_n(K)$. However, in our case we need a more powerful estimate that was proven by Vritsiou [168] in the context of showing the existence of regular M-ellipsoids for non-symmetric convex bodies.

Proposition 187 ([168], Corollary 11)**.** *Let $K \subseteq \mathbb{R}^n$ be a convex body so that $b(K) = \mathbf{0}$ and let $F \subseteq \mathbb{R}^n$ be a d-dimensional subspace. Then*

$$\mathrm{Vol}_d(\Pi_F(K))^{1/d} \lesssim \left(\frac{n}{d}\right)^5 \cdot \log\left(\frac{en}{d}\right)^2 \cdot \mathrm{Vol}_d(\Pi_F(K \cap -K))^{1/d}.$$

On a previous preprint, we had shown an inequality with better exponent when the body is centered so that $b(K^\circ) = \mathbf{0}$, i.e. the origin is the *Santaló point* of K. However, algorithmically the barycenter is much easier to compute and the exponent only affects the implicit universal constant in our main result, hence we choose to work with Vritsiou's estimate. For the interested reader, the bound with the Santaló point as center can be found in v2 on arXiv and also independently in [168].

We prove a custom-tailored inequality for later:

Lemma 188. *Let $K \subseteq \mathbb{R}^n$ be a convex body with $b(K) = \mathbf{0}$ and let $F \subseteq \mathbb{R}^n$ be a d-dimensional subspace. Then*

$$(\mathrm{Vol}_d(\Pi_F(K)))^{1/d} \lesssim \left(\frac{n}{d}\right)^6 \cdot \frac{\ell_{(K\cap -K)^\circ}}{d}.$$

Proof. We abbreviate $K_{\text{sym}} := K \cap -K$. Using the volume estimate from Proposition 187 with the assumption that the barycenter of K lies at the origin, we obtain

$$\begin{aligned}
(\mathrm{Vol}_d(\Pi_F(K)))^{1/d} &\overset{\text{Prop 187}}{\lesssim} \left(\frac{n}{d}\right)^6 \cdot (\mathrm{Vol}_d(\Pi_F(K_{\text{sym}})))^{1/d} \\
&\overset{\text{Cor 185}}{\lesssim} \left(\frac{n}{d}\right)^6 \cdot \frac{\ell_{(\Pi_F(K_{\text{sym}}))^\circ}}{d} \\
&\overset{\text{Lem 180}}{=} \left(\frac{n}{d}\right)^6 \cdot \frac{\ell_{K_{\text{sym}}^\circ \cap F}}{d} \\
&\overset{\text{Lem 183}}{\leq} \left(\frac{n}{d}\right)^6 \cdot \frac{\ell_{K_{\text{sym}}^\circ}}{d}.
\end{aligned}$$

Here we also used the fact that $(\Pi_F(K_{\text{sym}}))^\circ = K_{\text{sym}}^\circ \cap F$. $\square$

10.2.4 Properties of the covering radius

While the set K may not be symmetric, the sets Λ and $\mathbb{R}^n$ are symmetric, which implies the following:

Lemma 189 (Properties of the covering radius)**.** *Consider a lattice $\Lambda \subseteq \mathbb{R}^n$ and a compact convex set $K \subseteq \mathbb{R}^n$ with span(Λ) = affine.hull(K). Then*

(a) $\mu(\Lambda, K) = \mu(\Lambda, K + u)$ *for all* $u \in \mathrm{span}(\Lambda)$.

(b) $\mu(\Lambda, K) = \min\{r \geq 0 \mid (x + rK) \cap \Lambda \neq \emptyset \; \forall x \in \mathrm{span}(\Lambda)\}$.

We need a triangle inequality for the covering radius:

Lemma 190. *Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice and let $\Lambda' \subseteq \Lambda$ be a primitive sublattice. Then for any compact convex set $K \subseteq \mathbb{R}^n$ with $\mathbf{0} \in$ rel.int(K) and span(Λ) = span(K) one has*

$$\mu(\Lambda, K) \leq \mu(\Lambda', K \cap W) + \mu(\Lambda/\Lambda', \Pi_{W^\perp}(K)),$$

where $W :=$ span(Λ').

Proof. W.l.o.g. we may assume that Λ has full rank, so $\mathbf{0} \in \mathrm{int}(K)$. Following the characterization in Lemma 189.(b), we fix an $x \in \mathbb{R}^n$. For $r_1 := \mu(\Pi_{W^\perp}(\Lambda), \Pi_{W^\perp}(K))$ we know that $\Pi_{W^\perp}(x + r_1K) \cap \Pi_{W^\perp}(\Lambda) \neq \emptyset$. That means there is a $u_1 \in r_1K$ and a lattice point $y \in \Lambda$ so that $\Pi_{W^\perp}(x + u_1) = \Pi_{W^\perp}(y)$. Next, for $r_2 := \mu(\Lambda \cap W, K \cap W)$ we know that $(x + u_1 - y + r_2 \cdot (K \cap W)) \cap (\Lambda \cap W) \neq \emptyset$ which is equivalent to $(x + u_1 + r_2 \cdot (K \cap W)) \cap (y + (\Lambda \cap W)) \neq \emptyset$. Let $u_2 \in r_2 \cdot (K \cap W)$ be the vector so that $x + u_1 + u_2 \in \Lambda$. Then $u_1 + u_2 \in (r_1 + r_2)K$ by convexity, so $(x + (r_1 + r_2) \cdot K) \cap \Lambda \neq \emptyset$.

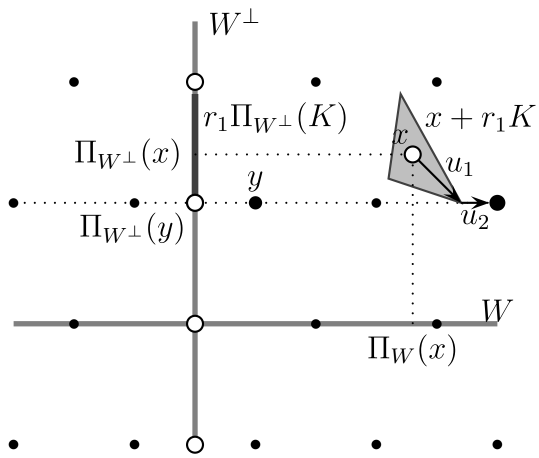

□

The natural extension of Lemma 190 to a filtration is as follows:

Lemma 191. *Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice with any sequence of sublattices $\{\mathbf{0}\} = \Lambda_0 \subset \Lambda_1 \subset \ldots \subset \Lambda_k = \Lambda$. Then for any compact convex set $K \subseteq \mathbb{R}^n$ with $\mathbf{0} \in rel.int(K)$ and $span(\Lambda) = span(K)$, one has*

$$\mu(\Lambda, K) \leq \sum_{i=1}^{k} \mu\left(\Lambda_i/\Lambda_{i-1}, \Pi_{span(\Lambda_{i-1})^\perp}(K \cap span(\Lambda_i))\right).$$

Proof. We can use the previous lemma to show by induction over $i_0 = k, k-1, \ldots, 1$ that

$$\mu(\Lambda, K) \leq \mu(\Lambda_{i_0-1}, K \cap \text{span}(\Lambda_{i_0-1})) + \sum_{i=i_0}^{k} \mu\left(\Lambda_i/\Lambda_{i-1}, \Pi_{\text{span}(\Lambda_{i-1})^\perp}(K \cap \text{span}(\Lambda_i))\right).$$

Indeed, for $i_0 = k$ this is exactly Lemma 190. If it holds for some $i_0 > 1$, then

$$\begin{aligned}\mu(\Lambda_{i_0-1}, K \cap \text{span}(\Lambda_{i_0-1})) \quad \leq \quad & \mu(\Lambda_{i_0-2}, K \cap \text{span}(\Lambda_{i_0-2})) + \\ & \mu\left(\Lambda_{i_0-1}/\Lambda_{i_0-2}, \Pi_{\text{span}(\Lambda_{i_0-2})^\perp}(K \cap \text{span}(\Lambda_{i_0-1}))\right),\end{aligned}$$

since $\text{span}(\Lambda_{i_0-2}) \subset \text{span}(\Lambda_{i_0-1})$. So the claim follows by induction, and taking $i_0 := 1$ yields the statement. □

10.2.5 Properties of μ_{KL}

We also need the following fact:

Lemma 192. *For any lattice $\Lambda \subseteq \mathbb{R}^n$, compact convex set K with $span(\Lambda) = affine.hull(K)$ and subspace $V \subseteq span(\Lambda)$ one has $\mu_{KL}(\Pi_V(\Lambda), \Pi_V(K)) \leq \mu_{KL}(\Lambda, K)$.*

Proof. Let $W \subseteq V$ be the subspace attaining the left side with $\dim W = d$. Then

$$\mu_{KL}(\Pi_V(\Lambda), \Pi_V(K)) = \left(\frac{\det(\Pi_W(\Pi_V(\Lambda)))}{\text{Vol}_d(\Pi_W(\Pi_V(K)))}\right)^{1/d} = \left(\frac{\det(\Pi_W(\Lambda))}{\text{Vol}_d(\Pi_W(K))}\right)^{1/d} \leq \mu_{KL}(\Lambda, K)$$

using that $\Pi_W(\Pi_V(x)) = \Pi_W(x)$ for all $x \in \mathbb{R}^n$ as $W \subseteq V$. □

10.2.6 Approximate stable lattices and the covering radius

Using the Reverse Minkowski Theorem it would not be hard to prove that for any stable lattice $\Lambda \subseteq \mathbb{R}^n$ one has $\mu(\Lambda, B_2^n) \leq O(\sqrt{n}\log(2n))$. In this section, we show how to generalize this to t-stable lattices and to general symmetric convex bodies. For a symmetric convex body Q, we consider the following quantity

$$\beta(Q) = \sup_{\Lambda \subseteq \mathbb{R}^n \text{ lattice}} \sup_{u \in \mathbb{R}^n} \frac{\rho_1((u+\Lambda) \setminus Q)}{\rho_1(\Lambda)}$$

Note that always $0 < \beta(Q) \leq 1$. Intuitively, a body Q with $\beta(Q) \ll 1$ is large enough that for any lattice a substantial fraction of the discrete Gaussian weight has to fall in Q. As part of the celebrated Transference Theorem, Banaszczyk showed how to relate the ℓ-value of a body to its β-value:

Lemma 193 (Banaszczyk [12])**.** *For any $\varepsilon > 0$, there is a $\delta > 0$ so that the following holds: for any symmetric convex body $Q \subseteq \mathbb{R}^n$ with $\ell_Q \leq \delta$ one has $\beta(Q) \leq \varepsilon$.*

Next, we can get a fairly tight upper bound on the covering radius of a t-stable lattice:

Proposition 194. *Let $\Lambda \subseteq \mathbb{R}^n$ be a full rank lattice that is the r-scaling of a t-stable lattice and let $Q \subseteq \mathbb{R}^n$ be a symmetric convex body. Then $\mu(\Lambda, Q) \leq O(\log(2n)) \cdot t \cdot r \cdot \ell_Q$.*

Proof. Let $\varepsilon > 0$ be a small enough constant that we determine later. Let δ be the constant so that Lemma 193 applies (w.r.t. ε). The claim is invariant under scaling Q, hence we may scale Q so that $\ell_Q \leq \delta$ and consequently $\beta(Q) \leq \varepsilon$. We may also scale the lattice so that Λ is t-stable (i.e. $r = 1$). It suffices to prove that under these assumptions, $\mu(\Lambda, Q) \leq s \cdot t$ where $s := C\log(2n)$ is the parameter from Lemma 179. Now suppose for the sake of contradiction that there is a translate $u \in \mathbb{R}^n$ so that $(u+\Lambda) \cap stQ = \emptyset$. Since $\beta(Q) \leq \varepsilon$, we know that

$$\rho_1\Big(\Big(\frac{u}{st} + \frac{\Lambda}{st}\Big) \setminus Q\Big) \leq \varepsilon \rho_1\Big(\frac{\Lambda}{st}\Big).$$

Multiplying the sets and parameters by st gives

$$\rho_{st}((u+\Lambda) \setminus stQ) \leq \varepsilon \rho_{st}(\Lambda). \quad (*)$$

Using that Λ is t-stable, we get

$$\frac{1}{3}\rho_{st}(\Lambda) \overset{\text{Lem 179}}{\leq} \rho_{st}(u+\Lambda) \overset{(u+\Lambda)\cap stQ=\emptyset}{=} \rho_{st}((u+\Lambda)\setminus stQ) \overset{(*)}{\leq} \varepsilon\rho_{st}(\Lambda).$$

Then choosing $\varepsilon \in (0, \frac{1}{3})$ gives a contradiction. $\square$

10.3 Overview

Goal of this section is to provide the reader with an overview and some intuition concerning the proof of our main result, Theorem 166. First, we want to prove the inequality from Theorem 166 (with an even better exponent) in the special case that both the lattice and the body K are well-scaled. We will not actually use Prop 195 later in this form, but it will provide us with the idea for a general proof strategy.

Proposition 195. *Let $\Lambda \subseteq \mathbb{R}^n$ be a full rank 2-stable lattice and let K be a convex body with $b(K) = \mathbf{0}$ so that $K \cap -K$ is in ℓ-position. Then $\mu(\Lambda, K) \leq O(\log^2(2n)) \cdot \mu_{KL}(\Lambda, K)$.*

Proof. We denote the inner symmetrizer by $K_{\text{sym}} := K \cap -K$. Then applying the estimate for stable lattices from Prop 194 we can upper bound the covering radius:

$$\mu(\Lambda, K) \overset{K\supseteq K_{\text{sym}}}{\leq} \mu(\Lambda, K_{\text{sym}}) \overset{\text{Prop 194}}{\lesssim} \log(2n)\cdot \ell_{K_{\text{sym}}}$$

Next, we lower bound $\mu_{KL}(\Lambda, K)$ by simply choosing the subspace $W := \mathbb{R}^n$ as witness. Then

$$\mu_{KL}(\Lambda, K) \geq \Big(\frac{\det(\Lambda)}{\mathrm{Vol}_n(K)}\Big)^{1/n} \overset{(*)}{\gtrsim} \frac{1}{\mathrm{Vol}_n(K_{\text{sym}})^{1/n}} \overset{\text{Cor 185}}{\gtrsim} \frac{n}{\ell_{K^\circ_{\text{sym}}}} \overset{\ell\text{-position}}{\gtrsim} \frac{\ell_{K_{\text{sym}}}}{\log(2n)}$$

where we use in $(*)$ that $\det(\Lambda) \geq 2^{-n}$ and $\mathrm{Vol}_n(K_{\text{sym}}) \geq 2^{-n}\mathrm{Vol}_n(K)$. Combining both inequalities gives the claim. $\square$

Next, we want to develop a proof strategy that works for general Λ and K. Translating K and applying a linear transformation to both Λ and K does not affect the claim, hence we may assume that K has the barycenter at $\mathbf{0}$ and the symmetrizer $K_{\text{sym}} := K \cap -K$ is in ℓ-position. But in general, Λ will not be a 2-stable lattice and we cannot expect that one can always choose the subspace $W = \mathbb{R}^n$ as witness like in Prop 195.

But we know by Cor 178 that the lattice Λ admits a 2-stable well-separated filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$. Let us abbreviate $d_i := \mathrm{rank}(\Lambda_i/\Lambda_{i-1})$ and $r_i := \det(\Lambda_i/\Lambda_{i-1})^{1/d_i}$. Then each quotient lattice $\frac{1}{r_i}\Lambda_i/\Lambda_{i-1}$ is a 2-stable lattice of dimension d_i and hence an argument similar to Prop 195 becomes feasible.

We can use the triangle inequality that we developed in Lemma 191 to obtain

$$\mu(\Lambda, K) \overset{K \supseteq K_{\text{sym}}}{\leq} \mu(\Lambda, K_{\text{sym}}) \overset{\text{Lem 191}}{\leq} \sum_{i=1}^{k} \mu(\Lambda_i/\Lambda_{i-1}, K_i) \overset{\text{Prop 194}}{\lesssim} \log(2n) \sum_{i=1}^{k} r_i \ell_{K_i} \leq \log(2n) \cdot r_k \ell_K$$

where $K_i := \Pi_{\text{span}(\Lambda_{i-1})^\perp}(K_{\text{sym}} \cap \text{span}(\Lambda_i))$. Here we have used that the sequence $r_1 < \ldots < r_k$ is geometrically increasing. This provides a convenient upper bound on the covering radius in terms of the relative determinant of the last quotient lattice in the filtration (which is the sparsest one). However we cannot avoid wondering whether we gave up too much by bounding $\ell_{K_i} \leq \ell_K$.

Next, we want to lower bound $\mu_{KL}(\Lambda, K)$. The only natural choices for a witness subspace seem to come from the filtration. Hence for some index $i \in \{1, \ldots, k\}$ we want to understand what can be obtained by choosing $W := \text{span}(\Lambda_{i-1})^\perp$, meaning we project out the densest $i-1$ of the quotient lattices. Then abbreviating $d := \dim(W) = d_i + \ldots + d_k$ we have

$$\mu_{KL}(\Lambda, K) \geq \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d} \overset{(*)}{\gtrsim} r_i \cdot \Big(\frac{d}{n}\Big)^6 \frac{d}{\ell^\circ_{K_{\text{sym}}}} \overset{\ell\text{-position}}{\gtrsim} r_i \cdot \log(2n) \cdot \Big(\frac{d}{n}\Big)^7 \cdot \ell_{K_{\text{sym}}}$$

In $(*)$ we use that $\Pi_W(\Lambda) = \Lambda/\Lambda_{i-1}$ and so $\det(\Pi_W(\Lambda))^{1/d}$ is a geometric mean of factors that are all at least r_i. Here we also use Lemma 188 to bound $\mathrm{Vol}_d(\Pi_W(K))$. It seems the only direct comparison can be obtained when letting $i := k$ in which case we have

$$\mu(\Lambda, K) \lesssim \log^2(2n) \cdot \Big(\frac{n}{d_k}\Big)^7 \cdot \mu_{KL}(\Lambda, K)$$

Hence, we can conclude Theorem 166 if d_k is close n, i.e. the last quotient subspace is large. But of course this is not necessarily true. In fact, the issue is more substantial. If K_{sym} is in ℓ-position with $\ell_{K_{\text{sym}}}$ and $\ell_{K_{\text{sym}}^{\circ}}$ known and W is a d-dimensional subspace, then this determines $\text{Vol}_d(\Pi_W(K))^{1/d}$ only up to a polynomial factor in $\frac{n}{d}$. Hence the information that we considered so far is simply too weak to approximate $\mu(\Lambda, K)$ up to a polylogarithmic factor. But fortunately there is a fix: instead of upper bounding the whole covering radius $\mu(\Lambda, K)$, we only estimate the covering radius corresponding to the less important half of the filtration. This means we will need to iterate the argument, which comes at the expense of a another logarithmic factor, but it will work!

10.4 *Proof of the main theorem*

We will spend the next two subsections proving our main Theorem 166 by induction over n. At each step, we split the lattice Λ and the convex body K into a subspace section of dimension at least $n/2$ and a projection where most of the work will go into analyzing the subspace section.

10.4.1 *The inductive step*

First, we give a self-contained description of the inductive step, then later in Section 10.4.2 we describe the main part of the induction.

Proposition 196. *There is a universal constant $C_0 > 0$ so that the following holds: For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$ with $b(K) = \mathbf{0}$, there exists a primitive sublattice $\Lambda' \subseteq \Lambda$ with $\text{rank}(\Lambda') \geq n/2$ so that*

$$\mu\big(\Lambda', (K \cap -K) \cap \textit{span}(\Lambda')\big) \leq C_0 \log^2(2n) \cdot \mu_{KL}(\Lambda, K).$$

Proof. Set $K_{\text{sym}} := K \cap (-K)$. The claim is invariant under applying a linear transformation to K and Λ. Hence we may assume that K_{sym} is in ℓ-position, i.e. $\ell_{K_{\text{sym}}} \cdot \ell_{K_{\text{sym}}^{\circ}} \leq O(n \log(2n))$. Consider a well-separated 2-stable filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$ which exists by Cor 178. We will later choose the lattice Λ' from one of the lattices Λ_i in the filtration, but we postpone the choice for now. We define

$$d_i := \text{rank}(\Lambda_i/\Lambda_{i-1}) \quad \text{and} \quad r_i := \text{nd}(\Lambda_i/\Lambda_{i-1}) = \det(\Lambda_i/\Lambda_{i-1})^{1/d_i},$$

which are the rank and normalized determinants of the quotient lattices in the filtration. Recall that $r_1 < r_2 < \ldots < r_k$ with $r_i \leq \frac{1}{2} r_{i+2}$ for all i.

Claim I. *For any $i \in \{1, \ldots, k\}$ one has $\mu(\Lambda_i, K_{sym} \cap span(\Lambda_i)) \lesssim \log(2n) \cdot r_i \cdot \ell_{K_{sym}}$.*

Proof of Claim I. We abbreviate $K_j := \Pi_{\text{span}(\Lambda_{j-1})^\perp}(K_{\text{sym}} \cap \text{span}(\Lambda_j))$. Then K_j is convex and symmetric and $\frac{1}{r_j}(\Lambda_j/\Lambda_{j-1})$ is a 2-stable lattice. Hence we can bound the covering radii of the individual quotient lattices by

$$\mu\big(\Lambda_j/\Lambda_{j-1}, K_j\big) \overset{\text{Prop 194}}{\lesssim} \log(2n) \cdot r_j \cdot \ell_{K_j} \overset{\text{Lem 184}}{\leq} \log(2n) \cdot r_j \cdot \ell_{K_{\text{sym}}}. \tag{10.2}$$

Then using the triangle inequality for the covering radius we bound

$$\begin{aligned} \mu(\Lambda_i, K_{\text{sym}} \cap \text{span}(\Lambda_i)) &\overset{\text{Lem 191}}{\leq} \sum_{j=1}^{i} \mu\left(\Lambda_j/\Lambda_{j-1}, K_j\right) \\ &\overset{(10.2)}{\lesssim} \log(2n) \cdot \ell_{K_{\text{sym}}} \cdot \sum_{j=1}^{i} r_j \\ &\lesssim \log(2n) \cdot \ell_{K_{\text{sym}}} \cdot r_i, \end{aligned}$$

using in the last step that $r_1 < \ldots < r_i$ and $r_j \leq \frac{1}{2} r_{j+2}$ for all j. □

In the following we abbreviate $d_{\geq i} := \sum_{j=i}^{k} d_j$.

Claim II. *For any $i \in \{1, \ldots, k\}$ one has $\mu_{KL}(\Lambda, K) \gtrsim \frac{r_i}{\log(2n)} \cdot (\frac{d_{\geq i}}{n})^7 \cdot \ell_{K_{sym}}$.*

Proof of Claim II. We choose the subspace $W := \text{span}(\Lambda_{i-1})^\perp$ as witness and note that $\Pi_W(\Lambda) = \Lambda/\Lambda_{i-1}$. Abbreviating $d := \dim(W) = \text{rank}(\Lambda/\Lambda_{i-1}) = d_{\geq i}$ we have

$$\det(\Lambda/\Lambda_{i-1})^{1/d} = \Big(\prod_{j=i}^{k} r_j^{d_j}\Big)^{1/\sum_{j=i}^{k} d_j} \geq r_i, \tag{10.3}$$

where the middle expression denotes a geometric average of values $r_i < r_{i+1} < \ldots < r_k$. Then lower bounding the covering radius proxy with the witness W gives

$$\begin{aligned} \mu_{KL}(\Lambda, K) &\geq \Big(\frac{\det(\Pi_W(\Lambda))}{\text{Vol}_d(\Pi_W(K))}\Big)^{1/d} \\ &\overset{(10.3)}{\geq} \frac{r_i}{\text{Vol}_d(\Pi_W(K))^{1/d}} \\ &\overset{\text{Lem 188}}{\gtrsim} r_i \cdot \Big(\frac{d}{n}\Big)^6 \cdot \frac{d}{\ell_{K_{\text{sym}}^\circ}} \overset{\ell\text{-position}}{\gtrsim} \frac{r_i}{\log(2n)} \cdot \Big(\frac{d}{n}\Big)^7 \cdot \ell_{K_{\text{sym}}}, \end{aligned}$$

using $\ell_{K_{\text{sym}}} \cdot \ell_{K^\circ_{\text{sym}}} \lesssim n\log(2n)$ in the last step. $\square$

Combining Claim I and Claim II with the same index i gives

$$\mu(\Lambda_i, K_{\text{sym}} \cap \text{span}(\Lambda_i)) \lesssim \log^2(2n) \cdot \Big(\frac{n}{d_{\geq i}}\Big)^7 \cdot \mu_{KL}(\Lambda, K).$$

Now, let $i^* \in \{1, \dots, k\}$ be the minimal index so that $\text{rank}(\Lambda_{i^*}) \geq \frac{n}{2}$. Then $d_{\geq i^*} \geq \frac{n}{2}$ by minimality. Hence $\Lambda' := \Lambda_{i^*}$ satisfies the claim. $\square$

10.4.2 *Completing the main proof*

Using Proposition 196 we can finish the proof of our main theorem.

Proof of Theorem 166. Consider a full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and a convex body $K \subseteq \mathbb{R}^n$. We will prove by induction over n that

$$\mu(\Lambda, K) \leq C_0 \log^3(2n) \cdot \mu_{KL}(\Lambda, K),$$

where $C_0 \geq 1$ is the constant from Proposition 196. The claim is true for $n = 1$, hence assume $n \geq 2$ from now on. The claim is invariant under translations of K, hence we may assume that $b(K) = \mathbf{0}$. Let $\Lambda' \subseteq \Lambda$ be the primitive sublattice from Prop 196 and set $W := \text{span}(\Lambda')$. Then

$$\begin{aligned}
\mu(\Lambda, K) &\stackrel{\text{Lem 190}}{\leq} \mu(\Lambda \cap W, K \cap W) + \mu(\Pi_{W^\perp}(\Lambda), \Pi_{W^\perp}(K)) \\
&\stackrel{K \supseteq K_{\text{sym}}}{\leq} \mu(\Lambda \cap W, K_{\text{sym}} \cap W) + \mu(\Pi_{W^\perp}(\Lambda), \Pi_{W^\perp}(K)) \\
&\stackrel{\substack{\text{Prop 196}\\ +\text{ind.}}}{\leq} C_0 \log^2(2n) \cdot \mu_{KL}(\Lambda, K) + C_0 \log^3(2\underbrace{\dim(W^\perp)}_{\leq n/2}) \cdot \underbrace{\mu_{KL}(\Pi_{W^\perp}(\Lambda), \Pi_{W^\perp}(K))}_{\leq \mu_{KL}(\Lambda, K)} \\
&\stackrel{\text{Lem 192}}{\leq} C_0 \underbrace{\log^2(2n) \cdot \Big(1 + \log(2n)\Big)}_{= \log^3(2n)} \cdot \mu_{KL}(\Lambda, K). \quad \square
\end{aligned}$$

We should point out that Regev and Stevens-Davidowitz [137] prove that in the Euclidean case one has $\mu(\Lambda, B_2^n) \leq O(\log^{3/2}(2n)) \cdot \mu_{KL}(\Lambda, B_2^n)$. Our proof could be seen as a generalization of their argument in the sense that [137] also relate both notions of covering radii to the quantities r_i and d_i as defined in Prop 196 by proving that

$$\mu(\Lambda, B_2^n) \leq O(\log(2n)) \cdot \sqrt{\sum_{i=1}^{k} d_i r_i^2} \leq O(\log^{3/2}(2n)) \cdot \mu_{KL}(\Lambda, B_2^n)$$

On the other hand, for them the "standard" canonical filtration suffices and they do not require an inductive step. Implicitly, our induction causes $O(\log(2n))$ many re-centering and rescaling operations using the result of Figiel, Tomczak-Jaegerman and Pisier (Theorem 182). This circumvents the issue that the covering radius might be dominated by a subspace of dimension d with $d \ll n$, which may not affect the ℓ-position of the body sufficiently. Then implicitly the induction will contain an iteration where d is relatively large compared to the current ambient dimension. It may also be instructive to reconsider the proof of Prop 196 in the case that $K = B_2^n$. Then in (10.2), we would obtain the inequality $\mu(\Lambda_j/\Lambda_{j-1}, K_j) \lesssim \log(2n) \cdot r_j \cdot \sqrt{n}$ while actually the much stronger bound of $\mu(\Lambda_j/\Lambda_{j-1}, K_j) \lesssim \log(2n) \cdot r_j \cdot \sqrt{d_j}$ holds. The trick is that using a well-separated filtration the arising loss can be efficiently bounded.

10.5 Finding the subspace W in single-exponential time

In this section, we prove Theorem 167, which guarantees that a suitable subspace subspace W can be found in time $2^{O(n)}$ at the expense of an additional logarithmic factor in the approximation guarantee. It will be convenient to first apply a linear transformation to well-scale K. This can be done in polynomial time and is a standard argument, see Lemma 203 for details. Hence, for us it suffices to prove the following:

Theorem 197. *Given a full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and a convex body $K \subseteq \mathbb{R}^n$ such that $B_2^n \subseteq K \subseteq (n+1)^{3/2} B_2^n$, there exists a randomized $2^{O(n)}$-time algorithm to compute a subspace $W \subseteq \mathbb{R}^n$ with $d := \dim(W)$ so that*

$$\mu(\Lambda, K) \lesssim \log^4(2n) \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d}.$$

The main technical tool will be the following result of Dadush, which is the only step in the algorithm which takes exponential time:

Theorem 198 (Theorem 6.4. in [49]). *Given a lattice $\Lambda \subseteq \mathbb{R}^n$ one can compute an $O(\log(2n))$-stable filtration of Λ in $2^{O(n)}$ time with probability at least $1 - 2^{-\Omega(n)}$.*

The following algorithm mimics the proof in Section 4:

FIND-SUBSPACE

Input: Convex body $K \subseteq \mathbb{R}^n$ so that $B_2^n \subseteq K \subseteq (n+1)^{3/2} B_2^n$, full rank $\Lambda \subseteq \mathbb{R}^n$

Output: Subspace $W \subseteq \mathbb{R}^n$ satisfying Theorem 197

(1) Compute an approximate barycenter $\tilde{x}$ such that $\|b(K) - \tilde{x}\|_2 \leq 1$

(2) Shift $K' := K - \tilde{x}$

(3) Set $K_{\text{sym}} := K' \cap (-K')$ and compute an invertible linear map T so that

$$\ell_{T(K_{\text{sym}})} \cdot \ell_{(T(K_{\text{sym}}))^\circ} \leq C \cdot n \log(2n)$$

(4) Set $K' \leftarrow T(K)$ and $\Lambda' \leftarrow T(\Lambda)$

(5) Compute an $O(\log(2n))$-stable filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda'$

(6) Compute a well-separated $O(\log(2n))$-stable filtration $\{\mathbf{0}\} = \Lambda'_0 \subset \ldots \subset \Lambda'_{k'} = \Lambda'$

(7) Set i^* as the minimal index with $\text{rank}(\Lambda'_{i^*}) \geq \frac{n}{2}$

(8) Set $W_{i^*} := \text{span}(\Lambda'_{i^*})^\perp$.

(9) Recursively call $W_\Pi :=$ FIND-SUBSPACE$(\Pi_{\text{span}(\Lambda'_{i^*})^\perp}(K')), \Pi_{\text{span}(\Lambda'_{i^*})^\perp}(\Lambda'))$

(10) Return $W := T^{-1} W'$ where

$$W' := \underset{W \in \{W_{i^*}, W_\Pi\}}{\text{argmin}} \Big\{ \Big(\frac{\det(\Pi_W(\Lambda'))}{\text{Vol}_{\dim(W)}(\Pi_W(K'))} \Big)^{1/\dim(W)} \Big\}.$$

We will need several volume computations in the algorithm, for which we use the following theorem:

Theorem 199 ([85]). *Given a convex body $K \subseteq \mathbb{R}^n$ with $r \cdot B_2^n \subseteq K \subseteq R \cdot B_2^n$, there exists a randomized algorithm which outputs a positive number ζ with $\mathrm{Vol}_n(K)/\zeta \in [1-\varepsilon, 1+\varepsilon]$. The runtime is polynomial in n, $1/\varepsilon$, $\log(1/r)$ and $\log(R)$.*

In fact, [85] also computes an approximation to the barycenter of K:

Theorem 200 ([85]). *Given a convex body $K \subseteq \mathbb{R}^n$ with $B_2^n \subseteq K \subseteq (n+1)^{3/2} \cdot B_2^n$ and $\delta > 0$, there exists a randomized algorithm with running time polynomial in n and $\frac{1}{\delta}$, which returns an approximate barycenter $\tilde{x}$ such that $\|b(K) - \tilde{x}\|_2 \leq \delta$.*

Now, we can prove the main result for this section:

Proof of Theorem 197. First we justify the running time of $2^{O(n)}$, later we discuss the approximation guarantee. We first apply Theorem 200 to compute an approximate barycenter $\tilde{x}$ and shift $K' := K - x$. Theorem 198 yields a filtration for step (5), which can be refined into a well-separated filtration by Theorem 177. Step (10) requires computation of determinants, which can be done in polynomial time via Gaussian elimination, and the volume of a convex body, which can also be done in randomized polynomial time using Theorem 199. The runtime $T(n)$ of Find-Subspace satisfies the recursion $T(n) \leq 2^{O(n)} + T(n/2)$, which can be resolved to $T(n) \leq 2^{O(n)}$.

Next, we justify the approximation guarantee. From the same argument in Section 3 and 4 one can see that the returned subspace satisfies

$$\mu(\Lambda, K) \lesssim \log^4(2n) \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d},$$

where we have taken into account that we pay an additional $\log(2n)$ factor from Proposition 194 as our filtration is only guaranteed to be $O(\log(2n))$-stable. Another subtle point is that we are using only an approximate barycenter. Hence it remains to generalize Proposition 187 and show that the approximation costs us at most another constant factor:

Claim. *Let $K \subseteq \mathbb{R}^n$ be a convex body so that $B_2^n \subseteq K$ and $\|b(K)\|_2 \leq 1$. Let $F \subseteq \mathbb{R}^n$ be a d-dimensional subspace. Then denoting $K_{\mathrm{sym}} := K \cap (-K)$,*

$$\mathrm{Vol}_d(\Pi_F(K))^{1/d} \lesssim \left(\frac{n}{d}\right)^5 \cdot \log\left(\frac{en}{d}\right)^2 \cdot \mathrm{Vol}_d(\Pi_F(K_{\mathrm{sym}}))^{1/d}.$$

Proof of Claim. By Proposition 187, we know that denoting $\tilde{K}_{\mathrm{sym}} := (K - b(K)) \cap (-K + b(K))$, we have

$$\mathrm{Vol}_d(\Pi_F(K))^{1/d} \lesssim \left(\frac{n}{d}\right)^5 \cdot \log\left(\frac{en}{d}\right)^2 \cdot \mathrm{Vol}_d(\Pi_F(\tilde{K}_{\mathrm{sym}}))^{1/d}.$$

Since $-b(K) \subseteq B_2^n \subseteq K$, it follows that $K - b(K) \subseteq K + K = 2K$, so that $\tilde{K}_{\mathrm{sym}} \subseteq 2K_{\mathrm{sym}}$ and $\mathrm{Vol}_d(\Pi_F(\tilde{K}_{\mathrm{sym}}))^{1/d} \leq 2 \cdot \mathrm{Vol}_d(\Pi_F(K_{\mathrm{sym}}))^{1/d}$. $\square$

10.6 Integer programming in time $(\log(2n))^{O(n)}$

Next, we show that integer programming can be solved in time $(\log(2n))^{O(n)}$. In fact, this is a known consequence of Theorem 167. We do not claim any original contribution for this section, but we reproduce the arguments of Dadush [58] to be self-contained. As it is common in the literature, we only state the dependence of running times on n; all running times that involve a convex set $K \subseteq rB_2^n$ and a lattice $\Lambda = \Lambda(B)$ also contain a not mentioned factor that is polynomial in $\log(r)$ and in the encoding length of B.

First, we describe the intuition behind Dadush's algorithm. Consider a convex body $K \subseteq \mathbb{R}^n$ and a lattice $\Lambda \subseteq \mathbb{R}^n$; the goal is to find a point in $K \cap \Lambda$. We compute a subspace $W \subseteq \mathbb{R}^n$ in time $2^{O(n)}$ that certifies the covering radius $\mu(\Lambda, K)$ up to a factor $\rho(n) := \Theta(\log^4(2n))$. Consider the points $X := \Pi_W(K) \cap \Pi_W(\Lambda)$ in the projection on W. For each $x \in K \cap \Lambda$, we also have $\Pi_W(x) \in X$. Note that the reverse may not be true in the sense that it is entirely possible that $K \cap \Lambda = \emptyset$ while $X \neq \emptyset$. However, we are guaranteed that all lattice points in K must be in one of the $(n-d)$-dimensional fibers of the projection, i.e.

$$K \cap \Lambda \subseteq \bigcup_{y \in X} \big((K \cap \Pi_W^{-1}(y)) \cap \Lambda\big).$$

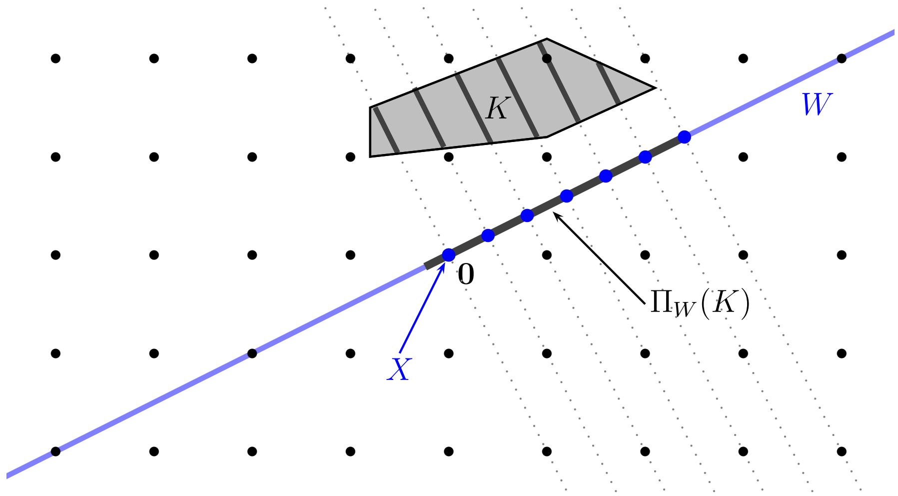

The algorithm enumerates X and then recurses on all the fibers. In order for this algorithm to be efficient we need to (i) bound the cardinality $|X|$ and (ii) be able to enumerate X. For (ii), note that it is possible that $W = \mathbb{R}^n$ and hence we would not gain anything by treating $\Pi_W(K) \cap \Pi_W(\Lambda)$ as a general integer programming problem.

For convex bodies $A, B \subseteq \mathbb{R}^n$, the *covering number* $N(A,B) := \min\{N \mid \exists x_1, \ldots, x_N \in \mathbb{R}^n : A \subseteq \bigcup_{i=1}^N (x_i + B)\}$ is the minimum number of translates of B needed to cover A. For a convex body $K \subseteq \mathbb{R}^n$ and a full rank lattice $\Lambda \subseteq \mathbb{R}^n$ we define

$$G(\Lambda, K) := \max_{x \in \mathbb{R}^n} |(K + x) \cap \Lambda|.$$

In words, $G(\Lambda, K)$ denotes the maximum number of lattice points that any shift of K contains. Note that even if $K \cap \Lambda = \emptyset$, $G(\Lambda, K)$ might still be arbitrarily large. However, algorithmically the quantity $G(\Lambda, K)$ is useful:

Theorem 201 ([54, 57])**.** *Given a convex body $K \subseteq \mathbb{R}^n$ and a full rank lattice $\Lambda \subseteq \mathbb{R}^n$, one can enumerate all points in $K \cap \Lambda$ in deterministic time $2^{O(n)} \cdot G(\Lambda, K)$.*

We briefly sketch the algorithm behind Theorem 201: We use the method of Dadush and Vempala [57] to compute an *M-ellipsoid* $\mathcal{E}$ of K which has the property that $N(K, \mathcal{E}), N(\mathcal{E}, K) \leq 2^{O(n)}$. Their deterministic algorithm takes time $2^{O(n)}$. In particular this means that $2^{-\Theta(n)} \leq$

$\frac{G(\Lambda,K)}{G(\Lambda,\mathcal{E})} \leq 2^{\Theta(n)}$. Next, we compute[1] the translates $x_1, \ldots, x_N$ with $N \leq 2^{O(n)}$ so that $K \subseteq \bigcup_{i=1}^N (x_i + \mathcal{E})$. Then we can use the following argument by Dadush, Peikert and Vempala [54] to enumerate all lattice points in $(x_i + \mathcal{E}) \cap \Lambda$. After applying a linear transformation, it suffices to compute all points in $(t + B_2^n) \cap \Lambda$ for $t \in \mathbb{R}^n$. Let $R \subseteq \Lambda \setminus \{\mathbf{0}\}$ be the *Voronoi-relevant* vectors, which are all the vectors that define a facet of the *Voronoi cell* of Λ. It is known that $|R| \leq 2^{n+1}$ and moreover the set R can be computed in time $2^{O(n)}$ by the algorithm of [117]. Next, consider the graph $H = (\Lambda, E)$ with edges $E = \{\{x,y\} : x, y \in \Lambda \text{ and } x - y \in R\}$. Then it follows from the work of [117] that the subgraph induced by $\Lambda \cap (t + B_2^n)$ is connected. Hence, one can compute the closest lattice point to t (again using [117]) and then traverse the subgraph.

Next, we require an upper bound on $G(\Lambda, K)$ in terms of the volume of K and density of Λ. Surprisingly, such an upper bound exists if we additionally control the covering radius. We reproduce Dadush's proof as the argument is key to understanding the algorithm:

Lemma 202. *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$ one has*

$$G(\Lambda, K) \leq 2^n \max\{\mu(\Lambda, K)^n, 1\} \cdot \frac{\mathrm{Vol}_n(K)}{\det(\Lambda)}.$$

Proof. After a linear transformation and scaling by $\max\{\mu(\Lambda, K), 1\}$, the statement is equivalent to the following simpler claim:

Claim. *For any convex body $K \subseteq \mathbb{R}^n$ with $\mu(\mathbb{Z}^n, K) \leq 1$ and any $x \in \mathbb{R}^n$ one has $|K \cap (x + \mathbb{Z}^n)| \leq 2^n \mathrm{Vol}_n(K)$.*

Proof of Claim. The claim is invariant under translating K, hence we may assume that $\mathbf{0} \in K$. Let $\equiv$ be the equivalence relation on pairs $x, y \in K$ that is defined by $x \equiv y \Leftrightarrow x - y \in \mathbb{Z}^n$. We define a set $V \subseteq K$ by selecting one element from each equivalence class w.r.t. $\equiv$. It would not matter much which element was selected, but let us make the canon-

[1] At least in the case that $\mathcal{E}$ is an M-ellipsoid for K, one may find those translates with $N \leq 2^{O(n)} N(K, \mathcal{E})$ with ease. After applying a linear transformation, we may assume that $\mathcal{E} = \sqrt{n} B_2^n$. Then take all translates $x + \mathcal{E}$ with $x \in \mathbb{Z}^n$ that intersect K.

ical choice of choosing the lexicographically minimal one. In other words, we choose

$$V = \{x \in K \mid x \leq_{\text{lex}} y \quad \forall y \in (x + \mathbb{Z}^n) \cap K\}$$

where $\leq_{\text{lex}}$ is the standard lexicographical ordering.

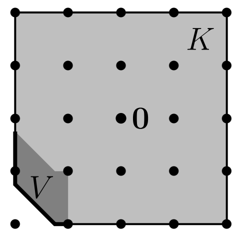

As we select at most one element from each equivalence class, we certainly have $\text{Vol}_n(V) \leq 1$. On the other hand, $\mu(\mathbb{Z}^n, K) \leq 1$ implies that for all $x \in \mathbb{R}^n$ one has $(x + \mathbb{Z}^n) \cap K \neq \emptyset$. That in turn means that every equivalence class has a member in K and so $\text{Vol}_n(V) \geq 1$. Together this gives $\text{Vol}_n(V) = 1$. Next, we note that by construction all translates $x + V$ with $x \in \mathbb{Z}^n$ are disjoint. Moreover, for $x \in K \cap \mathbb{Z}^n$ one has that $x + V \subseteq K + K = 2K$. Then

$$|K \cap \mathbb{Z}^n| = \sum_{x \in K \cap \mathbb{Z}^n} \underbrace{\text{Vol}_n(x + V)}_{=1} \stackrel{\text{disj.}}{=} \text{Vol}_n\Big(\bigcup_{x \in K \cap \mathbb{Z}^n} (x + V)\Big) \leq \text{Vol}_n(2K),$$

which gives the claim. □

One technicality we have to deal with is that Theorem 167 requires a lower bound on the *inradius* of K. Hence we run a preprocessing step: if there is no suitable lower bound for the inradius, then the lattice points of K are all contained in an easy-to-find hyperplane.

Lemma 203. *Given a compact convex set $K \subseteq rB_2^n$ and a lattice $\Lambda = \Lambda(B)$. Then in time polynomial in n, times a polynomial in $\log(r)$ and the encoding length of B one can find at least one of the following:*

(*a*) *An ellipsoid $\mathcal{E}$ and center c so that $c + \frac{1}{(n+1)^{3/2}}\mathcal{E} \subseteq K \subseteq c + \mathcal{E}$.*

(*b*) *A vector $a \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ and $\beta \in \mathbb{R}$ so that $K \cap \Lambda \subseteq \{x \in \mathbb{R}^n \mid \langle a, x\rangle = \beta\}$.*

Proof. We may assume that $\mathrm{rank}(\Lambda) = n$, otherwise any a orthogonal to $\mathrm{span}(\Lambda)$ will satisfy (b). Next, we use a variant of the ellipsoid method from [70] (see also Lemma 2.5.10 in [58]) to find a pair $(c, \mathcal{E})$ in time polynomial in n, $\log(r)$ and $\log(\frac{1}{\varepsilon})$ so that either (a) holds, or $K \subseteq c + \mathcal{E}$ and $\mathrm{Vol}_n(\mathcal{E}) \leq \varepsilon$. Suppose the latter happens. Then using Minkowski's Theorem (Theorem 174) in $(*)$ and the Blaschke-Santaló-Bourgain-Milman Theorem (Theorem 186) in $(**)$ we obtain

$$\lambda_1(\Lambda^*, \mathcal{E}^\circ) \stackrel{(*)}{\lesssim} \Big(\frac{\det(\Lambda^*)}{\mathrm{Vol}_n(\mathcal{E}^\circ)}\Big)^{1/n} \stackrel{(**)}{\lesssim} \Big(\frac{\mathrm{Vol}_n(\mathcal{E})}{\det(\Lambda) \cdot \nu_n^2}\Big)^{1/n} \lesssim n \cdot \Big(\frac{\varepsilon}{\det(\Lambda)}\Big)^{1/n} \leq \frac{1}{2} \cdot 2^{-n/2}$$

for a suitable choice of $\varepsilon > 0$. Then the LLL-algorithm [98] can find a dual lattice vector $a \in \Lambda^* \setminus \{\mathbf{0}\}$ with $\|a\|_{\mathcal{E}^\circ} \leq 2^{n/2} \cdot \lambda_1(\Lambda^*, \mathcal{E}^\circ) \leq \frac{1}{2}$. That vector a with $\beta := \lceil \langle a, c \rangle \rfloor$ will satisfy (b). $\square$

We are now ready to state the complete algorithm. As mentioned earlier, we denote $\rho(n) := \Theta(\log^4(2n))$ as the approximation factor from Theorem 167.

DADUSH'S ALGORITHM

Input: Compact convex set $K \subseteq \mathbb{R}^n$, lattice $\Lambda \subseteq \mathbb{R}^n$

Output: Point $x \in K \cap \Lambda$ or decision that there is none

(1) Use Lemma 203. If case (b) happens, obtain hyperplane H with $K \cap \Lambda \subseteq H$. Recurse on DADUSH$(K \cap H, \Lambda \cap H)$ and return the answer.

(2) Compute a subspace $W \subseteq \mathbb{R}^n$ with $d := \dim(W)$ and $R := (\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))})^{1/d}$ so that $R \leq \mu(\Lambda, K) \leq \rho(n) \cdot R$.

(3) Set $\tilde{K} := \min\{\rho(n) \cdot R, 1\} \cdot (K - c) + c$ for some $c \in K$.

(4) Compute an M-ellipsoid $\mathcal{E} \subseteq W$ for $\Pi_W(\tilde{K})$

(5) Compute $N \leq 2^{O(d)}$ points $x_1, \ldots, x_N \in W$ so $\Pi_W(\tilde{K}) \subseteq \bigcup_{i=1}^N (x_i + \mathcal{E})$.

(6) Compute $X := \Pi_W(\tilde{K}) \cap \Pi_W(\Lambda) = \big(\bigcup_{i=1}^N ((x_i + \mathcal{E}) \cap \Pi_W(\Lambda))\big) \cap \Pi_W(\tilde{K})$

(7) Recursively call DADUSH$(\tilde{K} \cap \Pi_W^{-1}(x), \Lambda \cap \Pi_W^{-1}(x))$ for all $x \in X$ and return any found lattice point (if there is any)

Here, to be more informative, we have expanded the blackbox from Theorem 201 into lines (4)-(6). The reader may also note a subtlety here that we have not discussed so far: if K

is very large so that $\mu(\Lambda, K) \ll 1$, then we may shrink K to a smaller body $\tilde{K} \subseteq K$ as long as we ensure that still $\mu(\Lambda, \tilde{K}) \leq 1$. We can now finish the analysis:

Theorem 204. *Dadush's algorithm finds a point in $K \cap \Lambda$ in time $(\log(2n))^{O(n)}$ if there is one.*

Proof. If the algorithm recurses in (1), the claim is clear by induction. So assume otherwise. First we argue correctness of the algorithm. Let $s := \min\{\rho(n) \cdot R, 1\} \in [0, 1]$ and recall that $\tilde{K} \subseteq K$ is a scaling of K by a factor of s. After step (3), the algorithm searches for a lattice point in $\tilde{K}$ rather than in the original body K. If $s < 1$, then the covering radius of the shrunk body is $\mu(\Lambda, \tilde{K}) = \frac{1}{\rho(n) \cdot R}\mu(\Lambda, K) \leq 1$. In other words, even though we continue the search in the strictly smaller body $\tilde{K}$, we are still guaranteed that $\tilde{K} \cap \Lambda \neq \emptyset$. Next, we discuss the running time of the algorithm. We estimate that

$$\begin{aligned}
G(\Pi_W(\Lambda), \Pi_W(\tilde{K})) &\overset{\text{Lem 202}}{\leq} 2^d \max\left\{\mu(\Pi_W(\Lambda), \Pi_W(\tilde{K}))^d, 1\right\} \cdot \frac{\mathrm{Vol}_d(\Pi_W(\tilde{K}))}{\det(\Pi_W(\Lambda))} \\
&\leq 2^d \max\left\{\Big(\underbrace{\frac{\rho(n)R}{s}}_{\geq 1}\Big)^d, 1\right\} \cdot s^d \cdot \underbrace{\frac{\mathrm{Vol}_d(\Pi_W(K))}{\det(\Pi_W(\Lambda))}}_{=R^{-d}} \\
&= 2^d \cdot (\rho(n)R)^d \cdot R^{-d} = (2\rho(n))^d.
\end{aligned}$$

Here we use that $\mu(\Pi_W(\Lambda), \Pi_W(\tilde{K})) \leq \mu(\Lambda, \tilde{K}) = \frac{1}{s} \cdot \mu(\Lambda, K) \leq \frac{\rho(n) \cdot R}{s}$. Then $|X| \leq G(\Pi_W(\Lambda), \Pi_W(\tilde{K})) \leq 2^d \rho(n)^d$ and by Lemma 202, the computation of X in (4)-(6) takes time $2^{O(d)}\rho(n)^d$. Now, let $T(n)$ be the maximum running time of the algorithm on n-dimensional instances. Then we have the recursion

$$T(n) \leq \max_{d \in \{1,\ldots,n\}} \left\{2^{O(n)} + (O(1) \cdot \rho(n))^d \cdot T(n-d)\right\} \quad \text{and} \quad T(1) = \Theta(1)$$

which indeed resolves to $T(n) \leq O(\rho(n))^n$. $\square$

We also explain how Dadush's algorithm can be used to solve integer linear programs in time $(\log(2n))^{O(n)}$. Again, the arguments used are standard. Details on the estimates can be found in the book of Schrijver [154].

Proof of Theorem 169. Consider an arbitrary integer linear program $\max\{c^T x \mid Ax \leq b, x \in \mathbb{Z}^n\}$. One can compute a number M in time polynomial in n and the encoding length of A and b so that if the IP is bounded and feasible, then the optimum value is the same as $\max\{c^T x \mid Ax \leq b, \|x\|_\infty \leq M, x \in \mathbb{Z}^n\}$. Next, by applying binary search, it suffices to find an integer point in the compact convex set $K = \{x \in \mathbb{R}^n \mid c^T x \geq \delta, Ax \leq b, \|x\|_\infty \leq M\}$ for which Theorem 168 applies. $\square$

10.7 Implications of Theorem 166

Here we derive a few implications of our main result. The following classical inequality will be useful here:

Lemma 205 ([143]). *For any convex set $K \subseteq \mathbb{R}^n$ we have $\mathrm{Vol}_n(K - K) \leq \binom{2n}{n} \cdot \mathrm{Vol}_n(K)$.*

We restate Theorem 170, which yields a nearly tight relationship between the covering radii of K and $K - K$. We remark that it remains an open question whether the two quantities are equal up to a constant.

Theorem (Theorem 170). *For any full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and any convex body $K \subseteq \mathbb{R}^n$, one has*

$$\mu(\Lambda, K - K) \leq \mu(\Lambda, K) \leq O(\log^3(2n)) \cdot \mu(\Lambda, K - K).$$

Proof. Let W denote the subspace attaining $\mu_{KL}(\Lambda, K)$ with $\dim W = d$. We can use Theorem 166 to upper bound

$$\begin{aligned} \mu(\Lambda, K) \lesssim \log^3(2n) \cdot \mu_{KL}(\Lambda, K) \quad &= \quad \log^3(2n) \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d} \\ &\overset{\text{Lem 205}}{\lesssim} \quad \log^3(2n) \cdot 4 \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K - K))}\Big)^{1/d} \\ &\lesssim \quad \log^3(2n) \cdot \mu_{KL}(\Lambda, K - K) \\ &\lesssim \quad \log^3(2n) \cdot \mu(\Lambda, K - K). \quad \square \end{aligned}$$

This in turn implies that the *flatness constant* in dimension n is bounded by $O(n \log^3(2n))$:

Theorem (Theorem 171)**.** *For any convex body $K \subseteq \mathbb{R}^n$ and any full rank lattice $\Lambda \subseteq \mathbb{R}^n$, one has*

$$\mu(\Lambda, K) \cdot \lambda_1(\Lambda^*, (K-K)^\circ) \leq O(n \log^3(2n)).$$

Proof. First we show a slightly worse bound of $O(n\log^4(2n))$. Banaszczyk [12] proved that for any symmetric convex body $Q \subseteq \mathbb{R}^n$ one has $\mu(\Lambda, Q) \cdot \lambda_1(\Lambda^*, Q^\circ) \leq O(n\log(2n))$. Setting $Q := K - K$ (which is a symmetric convex body) one then has by Theorem 170

$$\mu(\Lambda, K) \cdot \lambda_1(\Lambda^*, Q^\circ) \leq O(\log^3(2n)) \cdot \mu(\Lambda, Q) \cdot \lambda_1(\Lambda^*, Q^\circ) \leq O(n\log^4(2n)).$$

Now we give the argument of the stronger bound of $O(n\log^3(2n))$ which is due to Dadush. Let W denote the subspace attaining $\mu_{KL}(\Lambda, K)$ with $\dim W = d$. By Theorem 166,

$$\begin{aligned}
\mu(\Lambda, K) \lesssim \log^3(2n) \cdot \mu_{KL}(\Lambda, K) \quad &= \quad \log^3(2n) \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(K))}\Big)^{1/d} \\
&\overset{\text{Lem 205}}{\lesssim} \quad \log^3(2n) \cdot 4 \cdot \Big(\frac{\det(\Pi_W(\Lambda))}{\mathrm{Vol}_d(\Pi_W(Q))}\Big)^{1/d} \\
&\overset{\text{Lem 186}}{\asymp} \quad \log^3(2n) \cdot d \cdot \Big(\frac{\cdot \mathrm{Vol}_d(Q^\circ \cap W)}{\det(\Lambda^* \cap W)}\Big)^{1/d} \\
&\overset{\text{Thm 174}}{\lesssim} \quad n\log^3(2n) \cdot \frac{2}{\lambda_1(\Lambda^* \cap W, Q^\circ \cap W)}.
\end{aligned}$$

Here, we have used that $\Pi_W(\Lambda)^* = \Lambda^* \cap W$. Since $\lambda_1(\Lambda^*, Q^\circ) \leq \lambda_1(\Lambda^* \cap W, Q^\circ \cap W)$, the theorem follows. $\square$

We also explain the proof of Corollary 172 which again is standard:

Corollary (Cor 172)**.** *Let $K \subseteq \mathbb{R}^n$ by a convex body with $K \cap \mathbb{Z}^n = \emptyset$. Then there is a vector $c \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$ so that at most $O(n\log^3(2n))$ many hyperplanes of the form $\langle c, x\rangle = \delta$ with $\delta \in \mathbb{Z}$ intersect K.*

Proof. We apply Theorem 171 for the lattice $\Lambda := \mathbb{Z}^n$ so that $\Lambda^* = \mathbb{Z}^n$. Then $K \cap \mathbb{Z}^n = \emptyset$ implies that $\mu(\mathbb{Z}^n, K) > 1$ and so $\lambda_1(\mathbb{Z}^n, (K-K)^\circ) \lesssim n\log^3(2n)$. Let $c \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$ be the vector attaining this bound. Then revisiting the definition of the dual norm (Sec 10.2.3) we have $\max\{\langle c, x-y\rangle : x,y \in K\} = \|c\|_{(K-K)^\circ}$. That means at most $\|c\|_{(K-K)^\circ}+1 \lesssim n\log^3(2n)$ hyperplanes of the form $\langle c, x\rangle = \delta$ with $\delta \in \mathbb{Z}$ intersect K. $\square$

10.8 The approximate canonical filtration

In this chapter, we prove Theorem 177. The proof idea is rather simple: given a t-stable filtration $\{\mathbf{0}\} = \Lambda_0 \subset \ldots \subset \Lambda_k = \Lambda$, we select one index from every density class in order to make the filtration well-separated. But before we come to the main argument, we require two lemmas.

Lemma 206 (Grayson's parallelogram rule [43])**.** *For any two lattices $\Lambda, \Lambda' \subseteq \mathbb{R}^n$,*

$$\det(\Lambda)\cdot\det(\Lambda') \geq \det(\Lambda+\Lambda')\cdot\det(\Lambda\cap\Lambda').$$

A proof may also be found in Chapter 2 of [160]. The t-stable filtration can be used to obtain lower bounds on the determinant of any sublattice:

Lemma 207. *Let $\Lambda \subseteq \mathbb{R}^n$ be any lattice and let $\{\mathbf{0}\} = \Lambda_0 \subset \Lambda_1 \subset \ldots \subset \Lambda_k = \Lambda$ be a t-stable filtration. Then for any sublattice $\tilde{\Lambda} \subseteq \Lambda$ we have the inequality*

$$\mathrm{nd}(\tilde{\Lambda}) \geq t^{-1}\cdot\mathrm{nd}(\Lambda_1).$$

Proof. Let $r_i := \mathrm{nd}(\Lambda_i/\Lambda_{i-1}) = \det(\Lambda_i/\Lambda_{i-1})^{1/\mathrm{rank}(\Lambda_i/\Lambda_{i-1})}$ be the normalized determinant. We prove by induction on $i \in \{1,\ldots,k\}$ that the result holds for all lattices $\tilde{\Lambda} \subseteq \Lambda_i$. The base case follows as $\Lambda_1 = \Lambda_1/\Lambda_0$ is a scalar of the t-stable lattice $\frac{1}{\mathrm{nd}(\Lambda_1)}\Lambda_1$. Now suppose that $\tilde{\Lambda} \subseteq \Lambda_i$ for some $i > 1$. Note that $\Lambda_+ := \tilde{\Lambda} + \Lambda_{i-1}$ satisfies $\Lambda_{i-1} \subseteq \Lambda_+ \subseteq \Lambda_i$, so that $\Lambda_+/\Lambda_{i-1} \subseteq \Lambda_i/\Lambda_{i-1}$ and $\mathrm{nd}(\Lambda_+/\Lambda_{i-1}) \geq t^{-1}\cdot r_i > t^{-1}\cdot r_1$. By Lemma 206,

$$\det(\tilde{\Lambda})\cdot\det(\Lambda_{i-1}) \geq \det(\tilde{\Lambda}+\Lambda_{i-1})\cdot\det(\tilde{\Lambda}\cap\Lambda_{i-1}).$$

Factoring out Λ_{i-1} gives

$$\det(\tilde{\Lambda}) \geq \det(\Lambda_+/\Lambda_{i-1})\cdot\det(\tilde{\Lambda}\cap\Lambda_{i-1}).$$

Hence

$$\mathrm{nd}(\tilde{\Lambda}) \geq \mathrm{nd}(\Lambda_+/\Lambda_{i-1})^{\mathrm{rank}(\Lambda_+/\Lambda_{i-1})/\mathrm{rank}(\tilde{\Lambda})} \cdot \mathrm{nd}(\tilde{\Lambda} \cap \Lambda_{i-1})^{\mathrm{rank}(\tilde{\Lambda}\cap\Lambda_{i-1})/\mathrm{rank}(\tilde{\Lambda})} \geq t^{-1} \cdot r_1$$

where we used the inductive hypothesis on $\tilde{\Lambda} \cap \Lambda_{i-1} \subseteq \Lambda_{i-1}$ together with the fact that $\mathrm{rank}(\Lambda_+/\Lambda_{i-1}) + \mathrm{rank}(\tilde{\Lambda} \cap \Lambda_{i-1}) = \mathrm{rank}(\tilde{\Lambda})$. $\square$

Now, we come to the main argument:

Proof of Theorem 177. Let $r_i := \mathrm{nd}(\Lambda_i/\Lambda_{i-1})$ and $d_i :=$ rank$(\Lambda_i/\Lambda_{i-1})$. For $\ell \in \mathbb{Z}$ denote $I_\ell := \{i \in [k] : 2^\ell \leq r_i < 2 \cdot 2^\ell\}$. We define a sequence of indices $0 = \ell(0) < \ell(1) < \ldots < \ell(\tilde{k}) = k$ that contains precisely the largest index i in each I_ℓ with $I_\ell \neq \emptyset$ plus the index $\ell(0) = 0$. We set $\tilde{\Lambda}_j := \Lambda_{\ell(j)}$ and $\tilde{r}_j := \mathrm{nd}(\tilde{\Lambda}_j/\tilde{\Lambda}_{j-1})$. First, consider an index ℓ with $I_\ell \neq \emptyset$. Let $i_{\min}, i_{\max} \in I_\ell$ be the minimal and maximal indices in I_ℓ. Then

$$\begin{aligned}
\det(\Lambda_{i_{\max}}/\Lambda_{i_{\min}-1})^{1/\text{rank}(\Lambda_{i_{\max}}/\Lambda_{i_{\min}-1})} &= \Big(\prod_{i=i_{\min}}^{i_{\max}} \det(\Lambda_i/\Lambda_{i-1})\Big)^{1/\sum_{i=i_{\min}}^{i_{\max}} \text{rank}(\Lambda_i/\Lambda_{i-1})} \\
&= \Big(\prod_{i=i_{\min}}^{i_{\max}} r_i^{d_i}\Big)^{1/\sum_{i=i_{\min}}^{i_{\max}} d_i}.
\end{aligned}$$

Note that this value is a weighted geometric average of r_i-values for $i \in I_\ell$. From this it immediately follows that $\tilde{r}_1 < \ldots < \tilde{r}_k$ and $\tilde{r}_j \leq \frac{1}{2}\tilde{r}_{j+2}$ for all j, i.e. (a′) holds. It remains to show that the quotient lattices are scalars of $2t$-stable lattices. Fix some index $j \in [\tilde{k}]$ and let $\Lambda' := \frac{1}{\tilde{r}_j}(\tilde{\Lambda}_j/\tilde{\Lambda}_{j-1})$. First note that by assumption, the filtration $\{\mathbf{0}\} = \Lambda'_0 \subset \cdots \subset \Lambda'_{k'} := \Lambda'$ given by $\Lambda'_i := \frac{1}{\tilde{r}_j}(\Lambda_{\ell(j-1)+i}/\Lambda_{\ell(j-1)})$ with $k' := \ell(j) - \ell(j-1)$ is also t-stable because $\Lambda'_{i+1}/\Lambda'_i = \frac{1}{\tilde{r}_j}(\Lambda_{\ell(j-1)+i+1}/\Lambda_{\ell(j-1)+i})$.

We will prove the following two statements.

(I) For any sublattice $\tilde{\Lambda} \subseteq \Lambda'$ one has $\mathrm{nd}(\tilde{\Lambda}) \geq (2t)^{-1}$.

(II) For any sublattice $\tilde{\Lambda} \subseteq (\Lambda')^*$ one has $\mathrm{nd}(\tilde{\Lambda}) \geq (2t)^{-1}$.

First we show (I). We apply Lemma 207 on Λ' to obtain

$$\mathrm{nd}(\tilde{\Lambda}) \geq t^{-1} \cdot \mathrm{nd}(\Lambda'_1) \geq t^{-1} \cdot \frac{r_{\ell(j-1)+1}}{\tilde{r}_j} \geq (2t)^{-1},$$

since both numerator and denominator belong to the same interval $[2^\ell, 2{\cdot}2^\ell)$ for some $\ell \in \mathbb{Z}$. Next, we prove (II). Given the filtration $\{\mathbf{0}\} = \Lambda'_0 \subset \cdots \subset \Lambda'_{k'} = \Lambda'$ with $U_i := \mathrm{span}(\Lambda'_i)$, the dual filtration is given by $\{\mathbf{0}\} = (\Lambda')^*_0 \subset \cdots \subset (\Lambda')^*_{k'} = (\Lambda')^*$ with $(\Lambda')^*_i := \Lambda^* \cap U^\perp_{k'-i}$ and determinant $\det((\Lambda')^*_i) = \det((\Lambda')^*) \cdot \det(\Lambda'_{k'-i}) = \det(\Lambda'_{k'-i})$, see for example [49]. Since quotients of the dual filtration are duals of the quotients of the original filtration, the dual filtration is also t-stable. We then apply Lemma 207 on $(\Lambda')^*$:

$$\mathrm{nd}(\tilde{\Lambda}) \geq t^{-1} \cdot \mathrm{nd}((\Lambda')^*_1) = t^{-1} \cdot (r'_{k'})^{-1} = t^{-1} \cdot \left(\frac{r_{\ell(j)}}{\tilde{r}_j}\right)^{-1} \overset{r_{\ell(j)} \leq 2\cdot\tilde{r}_j}{\geq} (2t)^{-1}. \qquad \square$$

BIBLIOGRAPHY

[1] R. Ahlswede and A. Winter, *Strong converse for identification via quantum channels*, IEEE Transactions on Information Theory, 48 (2002), pp. 569–579.

[2] M. Ajtai, J. Aspnes, M. Naor, Y. Rabani, L. J. Schulman, and O. Waarts, *Fairness in scheduling*, Journal of Algorithms, 29 (1998), pp. 306–357.

[3] M. Ajtai, R. Kumar, and D. Sivakumar, *A sieve algorithm for the shortest lattice vector problem*, in Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece, 2001, pp. 601–610.

[4] N. Alon and J. H. Spencer, *The probabilistic method*, Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, fourth ed., 2016.

[5] R. Alweiss, Y. P. Liu, and M. Sawhney, *Discrepancy minimization via a self-balancing walk*, in Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2021, New York, NY, USA, 2021, Association for Computing Machinery, pp. 14–20.

[6] N. Anari and T.-D. Vuong, *An extension of Plücker relations with applications to subdeterminant maximization*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

[7] S. Artstein-Avidan, A. Giannopoulos, and V. D. Milman, *Asymptotic geometric analysis. Part I*, vol. 202 of Mathematical Surveys and Monographs, American Mathematical Society, Providence, RI, 2015.

[8] J. Aru, B. Narayanan, A. Scott, and R. Venkatesan, *Balancing sums of random vectors*, Discrete Analysis, 2018 (2018).

[9] K. Ball, E. A. Carlen, and E. H. Lieb, *Sharp uniform convexity and smoothness inequalities for trace norms*, Inventiones Mathematicae, 115 (1994), pp. 463–482.

[10] K. Ball and A. Pajor, *Convex bodies with few faces*, Proceedings of the American Mathematical Society, 110 (1990), pp. 225–231.

[11] W. Banaszczyk, *Balancing vectors and convex bodies*, Studia Mathematica, 106 (1993), pp. 93–100.

[12] W. Banaszczyk, *Inequalities for convex bodies and polar reciprocal lattices in $\mathbb{R}^n$ II: application of K-convexity*, Discret. Comput. Geom., 16 (1996), pp. 305–311.

[13] W. Banaszczyk, *Balancing vectors and Gaussian measures of n-dimensional convex bodies*, Random Structures Algorithms, 12 (1998), pp. 351–360.

[14] W. Banaszczyk, *On series of signed vectors and their rearrangements*, Random Structures & Algorithms, 40 (2012), pp. 301–316.

[15] W. Banaszczyk, A. E. Litvak, A. Pajor, and S. J. Szarek, *The flatness theorem for nonsymmetric convex bodies via the local theory of Banach spaces*, Math. Oper. Res., 24 (1999), pp. 728–750.

[16] N. Bansal, *Constructive algorithms for discrepancy minimization*, in 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA, 2010, pp. 3–10.

[17] N. Bansal, M. Charikar, R. Krishnaswamy, and S. Li, *Better algorithms and hardness for broadcast scheduling via a discrepancy approach*, in Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2014, pp. 55–71.

[18] 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), IEEE, 2016, pp. 788–799.

[19] 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, 2018, pp. 587–597.

[20] N. Bansal and S. Garg, *Algorithmic discrepancy beyond partial coloring*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, 2017, pp. 914–926.

[21] N. Bansal, H. Jiang, and R. Meka, *Resolving Matrix Spencer conjecture up to polylogarithmic rank*, 2022.

[22] 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), SIAM, 2021, pp. 2842–2861.

[23] N. Bansal, H. Jiang, R. Meka, S. Singla, and M. Sinha, *Prefix discrepancy, smoothed analysis, and combinatorial vector balancing*, in 13th Innovations in Theoretical Computer Science Conference, ITCS 2022, M. Braverman, ed., Leibniz International Proceedings in Informatics, LIPIcs, Germany, Jan. 2022, Schloss Dagstuhl- Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing.

[24] N. Bansal, H. Jiang, S. Singla, and M. Sinha, *Online vector balancing and geometric discrepancy*, CoRR, abs/1912.03350 (2019).

[25] N. Bansal and V. Nagarajan, *Approximation-friendly discrepancy rounding*, in A Journey Through Discrete Mathematics, Springer, 2017, pp. 89–114.

[26] N. Bansal, L. Rohwedder, and O. Svensson, *Flow time scheduling and prefix Beck-Fiala*, in Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022, New York, NY, USA, 2022, Association for Computing Machinery, pp. 331–342.

[27] N. Bansal and J. H. Spencer, *On-line balancing of random inputs*, Random Structures & Algorithms, 57 (2020), pp. 879–891.

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

[29] S. Barman, *Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory's theorem*, in Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC '15, New York, NY, USA, 2015, Association for Computing Machinery, pp. 361–369.

[30] F. Barthe, O. Guédon, S. Mendelson, and A. Naor, *A probabilistic approach to the geometry of the ℓ_p^n -ball*, Ann. Probab., 33 (2005), pp. 480–513.

[31] J. Batson, D. Spielman, and N. Srivastava, *Twice-Ramanujan sparsifiers*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009, 2009, pp. 255–262.

[32] J. Beck, *Roth's estimate of the discrepancy of integer sequences is nearly sharp*, Combinatorica, 1 (1981), pp. 319–325.

[33] J. Beck and T. Fiala, *"Integer-making" theorems*, Discrete Applied Mathematics, 3 (1981), pp. 1–8.

[34] V. Bhargava, S. Saraf, and I. Volkovich, *Deterministic factorization of sparse polynomials with bounded individual degree*, J. ACM, 67 (2020).

[35] A. A. Borovkov and K. A. Borovkov, *Asymptotic Analysis of Random Walks: Heavy-Tailed Distributions*, Encyclopedia of Mathematics and its Applications, Cambridge University Press, 2008.

[36] J. Bourgain, J. Lindenstrauss, and V. Milman, *Approximation of zonoids by zonotopes*, Acta Mathematica, 162 (1989), pp. 73–141.

[37] J. Bourgain and J. Nelson, *Toward a unified theory of sparse dimensionality reduction in Euclidean space*, CoRR, abs/1311.2542 (2013).

[38] R. Bozzai, V. Reis, and T. Rothvoss, *The vector balancing constant for zonotopes*, in 64th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2023, November 6-9, 2023, Santa Cruz, California, USA, 2023.

[39] S. Bubeck, *Convex optimization: Algorithms and complexity*, 2015.

[40] B. Bukh, *An improvement of the Beck-Fiala theorem*, Combinatorics, Probability and Computing, 25 (2016), pp. 380–398.

[41] I. Bárány and V. Grinberg, *On some combinatorial questions in finite-dimensional spaces*, Linear Algebra and its Applications, 41 (1981), pp. 1–9.

[42] E. A. Carlen, *A remainder term for Hölder's inequality for matrices and quantum entropy inequalities*, 2016.

[43] B. Casselman, *Stability of Lattices and the Partition of Arithmetic Quotients*, Asian Journal of Mathematics, 8 (2004), pp. 607–638.

[44] 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, SIAM, 2011, pp. 1607–1614.

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

[46] S. Chobanyan, *Convergence as of rearranged random series in Banach space and associated inequalities*, in Probability in Banach Spaces, 9, Springer, 1994, pp. 3–29.

[47] M. B. Cohen and R. Peng, *ℓ_p row sampling by Lewis weights*, in Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC '15, New York, NY, USA, 2015, Association for Computing Machinery, pp. 183–192.

[48] C. W. Combettes and S. Pokutta, *Revisiting the approximate Carathéodory problem via the Frank-Wolfe algorithm*, Mathematical Programming, (2021).

[49] D. Dadush, *On approximating the covering radius and finding dense lattice subspaces*, in Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, New York, NY, USA, 2019, Association for Computing Machinery, pp. 1021–1026.

[50] D. Dadush, F. Eisenbrand, and T. Rothvoss, *From approximate to exact integer programming*, in Integer Programming and Combinatorial Optimization: 24th International Conference, IPCO 2023, Madison, WI, USA, June 2123, 2023, Proceedings, Berlin, Heidelberg, 2023, Springer-Verlag, pp. 100–114.

[51] D. Dadush, S. Garg, S. Lovett, and A. Nikolov, *Towards a constructive version of Banaszczyk's vector balancing theorem*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016, September 7-9, 2016, Paris, France, 2016, pp. 28:1–28:12.

[52] D. Dadush, H. Jiang, and V. Reis, *A new framework for matrix discrepancy: Partial coloring bounds via mirror descent*, in Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022, June 20-24, 2022, Rome, Italy, Association for Computing Machinery, 2022, pp. 649–658.

[53] D. Dadush, A. Nikolov, K. Talwar, and N. Tomczak-Jaegermann, *Balancing vectors in any norm*, in 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), 2018, pp. 1–10.

[54] D. Dadush, C. Peikert, and S. Vempala, *Enumerative lattice algorithms in any norm via M-ellipsoid coverings*, in FOCS, IEEE Computer Society, 2011, pp. 580–589.

[55] D. Dadush and O. Regev, *Towards strong reverse Minkowski-type inequalities for lattices*, in FOCS, IEEE Computer Society, 2016, pp. 447–456.

[56] D. Dadush and S. Vempala, *Deterministic construction of an approximate M-ellipsoid and its applications to derandomizing lattice algorithms*, in Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12, USA, 2012, Society for Industrial and Applied Mathematics, pp. 1445–1456.

[57] D. Dadush and S. Vempala, *Near-optimal deterministic algorithms for volume computation via M-ellipsoids*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 19237–19245.

[58] D. N. Dadush, *Integer Programming, Lattice Algorithms, and Deterministic Volume Estimation*, PhD thesis, Georgia Institute of Technology, USA, 2012.

[59] A. Deligkas, J. Fearnley, T. Melissourgos, and P. G. Spirakis, *Approximating the existential theory of the reals*, Journal of Computer and System Sciences, 125 (2022), pp. 106–128.

[60] A. Dvoretzky, *Convexity, volume 7 of Proceedings of Symposia in Pure Mathematics*, vol. 7, Amer. Math. Soc., 1963.

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

[62] A. Fiat, A. R. Karlin, E. Koutsoupias, C. Mathieu, and R. Zach, *Carpooling in social networks*, in Extended Abstracts Summer 2015: Strategic Behavior in Combinatorial Structures; Quantitative Finance, Springer, 2017, pp. 29–34.

[63] T. Figiel and N. Tomczak-Jaegermann, *Projections onto Hilbertian subspaces of Banach spaces*, Israel Journal of Mathematics, 33 (1979), pp. 155–171.

[64] M. Fradelizi, *Sections of convex bodies through their centroid*, Arch. Math. (Basel), 69 (1997), pp. 515–522.

[65] C. Franks and M. Saks, *On the discrepancy of random matrices with many columns*, Random Structures & Algorithms, 57 (2020), pp. 64–96.

[66] D. A. Freedman, *On tail probabilities for martingales*, the Annals of Probability, 3 (1975), pp. 100–118.

[67] A. Giannopoulos, *On some vector balancing problems*, Studia Mathematica, 122 (1997), pp. 225–234.

[68] E. D. Gluskin, *Extremal properties of orthogonal parallelepipeds and their applications to the geometry of banach spaces*, Mathematics of the USSR-Sbornik, 64 (1989), p. 85.

[69] V. Grinberg and S. Sevastyanov, *Value of the Steinitz constant*, Functional Analysis and Its Applications, 14 (1980), pp. 125–126.

[70] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, vol. 2 of Algorithms and Combinatorics, Springer, 1988.

[71] C. Harshaw, F. Sävje, D. Spielman, and P. Zhang, *Balancing covariates in randomized experiments with the gram-schmidt walk design*, 2019.

[72] N. Harvey and S. Samadi, *Near-optimal herding*, in Annual Conference Computational Learning Theory, 2014.

[73] A. Hinrichs, J. Prochno, and J. Vybiral, *Entropy numbers of embeddings of Schatten classes*, Journal of Functional Analysis, 273 (2017), pp. 3241–3261.

[74] R. Hoberg and T. Rothvoss, *A logarithmic additive integrality gap for bin packing*, in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17, USA, 2017, Society for Industrial and Applied Mathematics, pp. 2616–2625.

[75] ——, *A Fourier-analytic approach for the discrepancy of random set systems*, in Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19, USA, 2019, Society for Industrial and Applied Mathematics, p. 25472556.

[76] S. B. Hopkins, P. Raghavendra, and A. Shetty, *Matrix discrepancy from quantum communication*, in STOC, Rome, Italy, 2022, Association for Computing Machinery, pp. 637–648.

[77] A. Jambulapati, V. Reis, and K. Tian, *Linear-sized sparsifiers via near-linear time discrepancy theory*, 2023.

[78] H. Jiang, J. Kulkarni, and S. Singla, *Online geometric discrepancy for stochastic arrivals with applications to envy minimization*, arXiv preprint arXiv:1910.01073, (2019).

[79] H. Jiang and V. Reis, *A tighter relation between hereditary discrepancy and determinant lower bound*, in SIAM Symposium on Simplicity in Algorithms, 2022.

[80] F. John, *Extremum problems with inequalities as subsidiary conditions*, in Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948, Interscience Publishers, Inc., New York, 1948, pp. 187–204.

[81] R. Kannan, *Improved algorithms for integer programming and related lattice problems*, in Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, STOC '83, New York, NY, USA, 1983, ACM, pp. 193–206.

[82] R. Kannan, *Minkowski's convex body theorem and integer programming*, Math. Oper. Res., 12 (1987), pp. 415–440.

[83] R. Kannan and L. Lovász, *Covering minima and lattice-point-free convex bodies*, Annals of Mathematics, 128 (1988), pp. 577–602.

[84] R. Kannan, L. Lovász, and M. Simonovits, *Isoperimetric problems for convex bodies and a localization lemma*, Discrete and computational geometry, 13 (1995), pp. 541–560.

[85] R. Kannan, L. M. Lovász, and M. Simonovits, *Random walks and an $O(n^5)$ volume algorithm for convex bodies*, Random Structures and Algorithms, 11 (1997), pp. 1–50.

[86] C. G. Khatri, *On certain inequalities for normal distributions and their applications to simultaneous confidence bounds*, Ann. Math. Statist., 38 (1967), pp. 1853–1867.

[87] J. H. Kim, J. Matoušek, and V. H. Vu, *Discrepancy after adding a single set*, Combinatorica, 25 (2005), p. 499.

[88] P. Klein and N. E. Young, *On the number of iterations for Dantzig–Wolfe optimization and packing-covering approximation algorithms*, SIAM Journal on Computing, 44 (2015), pp. 1154–1172.

[89] H. König and V. D. Milman, *On the covering numbers of convex bodies*, in Geometrical Aspects of Functional Analysis, Springer, 1987, pp. 82–95.

[90] J. Kulkarni, V. Reis, and T. Rothvoss, *Optimal online discrepancy minimization*, 2023.

[91] D. Kunisky, *The discrepancy of unsatisfiable matrices and a lower bound for the komlós conjecture constant*, SIAM Journal on Discrete Mathematics, 37 (2023), pp. 586–603.

[92] R. Latala, *Estimation of moments of sums of independent real random variables*, The Annals of Probability, 25 (1997), pp. 1502–1513.

[93] R. Latala and D. Matlak, *Royen's proof of the Gaussian correlation inequality*, Geometric Aspects of Functional Analysis, (2017), pp. 265–275.

[94] L. C. Lau, R. Ravi, and M. Singh, *Iterative methods in combinatorial optimization*, vol. 46, Cambridge University Press, 2011.

[95] M. Ledoux and M. Talagrand, *Probability in Banach spaces*, Classics in Mathematics, Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition.

[96] Y. T. Lee and H. Sun, *An SDP-based algorithm for linear-sized spectral sparsification*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, New York, NY, USA, 2017, Association for Computing Machinery, pp. 678–687.

[97] Y. T. Lee and M.-C. Yue, *Universal barrier is n-self-concordant*, Mathematics of Operations Research, 46 (2021), pp. 1129–1148.

[98] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann., 261 (1982), pp. 515–534.

[99] J. Lenstra, H. W., *Integer programming with a fixed number of variables*, Mathematics of Operations Research, 8 (1983), pp. 538–548.

[100] A. Levy, H. Ramadas, and T. Rothvoss, *Deterministic discrepancy minimization via the multiplicative weight update method*, in International Conference on Integer Programming and Combinatorial Optimization, Springer, 2017, pp. 380–391.

[101] L. Li and A. Nikolov, *On the gap between hereditary discrepancy and the determinant lower bound*, 2023.

[102] L. Lovász, J. Spencer, and K. Vesztergombi, *Discrepancy of set-systems and matrices*, Eur. J. Comb., 7 (1986), pp. 151–160.

[103] S. Lovett and R. Meka, *Constructive discrepancy minimization by walking on the edges*, in 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, 2012, pp. 61–67.

[104] F. Lust-Piquard and G. Pisier, *Non commutative Khintchine and Paley inequalities*, Arkiv för matematik, 29 (1991), pp. 241–260.

[105] P. Manurangsi and W. Suksompong, *Almost envy-freeness for groups: Improved bounds via discrepancy theory*, Theoretical Computer Science, 930 (2022), pp. 179–195.

[106] A. Marcus, D. Spielman, and N. Srivastava, *Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem*, Annals of Mathematics, 182 (2015), pp. 327–350.

[107] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families ii: Mixed characteristic polynomials and the Kadison-Singer problem*, Annals of Mathematics, 182 (2015), pp. 327–350.

[108] J. Matoušek, *The determinant bound for discrepancy is almost tight*, Proceedings of the American Mathematical Society, 141 (2013), pp. 451–460.

[109] J. Matoušek, A. Nikolov, and K. Talwar, *Factorization norms and hereditary discrepancy*, arXiv preprint arXiv:1408.1376, (2014).

[110] J. Matoušek, *An ℓ_p version of the Beck-Fiala conjecture*, Eur. J. Comb., (1998).

[111] J. Matoušek, *Geometric discrepancy: An illustrated guide*, vol. 18, Springer Science & Business Media, 2009.

[112] M. Meckes, *On the spectral norm of a random Toeplitz matrix*, Electronic Communications in Probability, 12 (2007), pp. 315–325.

[113] R. Meka, *Discrepancy and beating the union bound*, https://windowsontheory.org/2014/02/07/discrepancy-and-beating-the-union-bound/, 2014.

[114] M. Meyer and A. Pajor, *Sections of the unit ball of ℓ_p^n*, Journal of Functional Analysis, 80 (1988), pp. 109–123.

[115] M. Meyer and A. Pajor, *On the Blaschke-Santaló inequality*, Archiv der Mathematik, 55 (1990), pp. 82–93.

[116] M. Meyer and E. Werner, *The Santaló-regions of a convex body*, Transactions of the American Mathematical Society, 350 (1998), pp. 4569–4591.

[117] D. Micciancio and P. Voulgaris, *A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, 2010, pp. 351–358.

[118] V. S. Mirrokni, R. P. Leme, A. Vladu, and S. C. Wong, *Tight bounds for approximate Carathéodory and beyond*, in Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, D. Precup and Y. W. Teh, eds., vol. 70 of Proceedings of Machine Learning Research, PMLR, 2017, pp. 2440–2448.

[119] S. Myroshnychenko, M. Stephen, and N. Zhang, *Grünbaum's inequality for sections*, Journal of Functional Analysis, 275 (2018), pp. 2516–2537.

[120] M. Naor, *On fairness in the carpool problem*, Journal of Algorithms, 55 (2005), pp. 93–98.

[121] A. S. Nemirovski and D. B. Yudin, *Problem complexity and method efficiency in optimization*, Wiley-Interscience, 1983.

[122] Y. Nesterov and A. Nemirovski, *Interior-point polynomial algorithms in convex programming*, in Siam studies in applied mathematics, 1994.

[123] A. Newman, O. Neiman, and A. Nikolov, *Beck's three permutations conjecture: A counterexample and some consequences*, in 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, IEEE, 2012, pp. 253–262.

[124] A. Nikolov, *Randomized rounding for the largest simplex problem*, in Proceedings of the forty-seventh annual ACM symposium on Theory of computing, 2015, pp. 861–870.

[125] A. Nikolov, *Personal communication*, 2019.

[126] A. Nikolov, K. Talwar, and L. Zhang, *The geometry of differential privacy: the sparse and approximate cases*, in STOC, ACM, 2013, pp. 351–360.

[127] A. M. Odlyzko, *The rise and fall of knapsack cryptosystems*, Cryptology and Computational Number Theory, (1990), pp. 75–88.

[128] D. Pálvölgyi, *Indecomposable coverings with concave polygons*, Discrete & Computational Geometry, 44 (2010), pp. 577–588.

[129] C. Peikert, *Lattices in cryptography*, https://web.eecs.umich.edu/~cpeikert/lic15/, 2015.

[130] G. Pisier, *Sur les espaces de banach k-convexes*, Séminaire Analyse fonctionnelle (dit" Maurey-Schwartz"), (1979), pp. 1–15.

[131] G. Pisier, *Un théorème sur les opérateurs linéaires entre espaces de banach qui se factorisent par un espace de hilbert*, in Annales scientifiques de l'École Normale Supérieure, vol. 13, 1980, pp. 23–43.

[132] G. Pisier, *Remarques sur un résultat non publié de B. Maurey*, Séminaire d'Analyse fonctionnelle (dit "Maurey-Schwartz"), (1980-1981).

[133] G. Pisier, *The Volume of Convex Bodies and Banach Space Geometry*, Cambridge Tracts in Mathematics, Cambridge University Press, 1989.

[134] G. Pisier, *Introduction to operator space theory*, vol. 294, Cambridge University Press, 2003.

[135] O. Regev, *Lecture notes on lattices*, `http://www.cims.nyu.edu/~regev/teaching/lattices_fall_2009`, 2009.

[136] O. Regev, *On lattices, learning with errors, random linear codes, and cryptography*, J. ACM, 56 (2009), pp. 34:1–34:40.

[137] O. Regev and N. Stephens-Davidowitz, *A reverse Minkowski theorem*, in STOC, ACM, 2017, pp. 941–953.

[138] V. Reis and T. Rothvoss, *Linear size sparsifier and the geometry of the operator norm ball*, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, SIAM, 2020, pp. 2337–2348.

[139] V. Reis and T. Rothvoss, *Approximate Carathéodory bounds via discrepancy theory*, 2022.

[140] V. Reis and T. Rothvoss, *Vector balancing in Lebesgue spaces*, Random Structures and Algorithms, (2022).

[141] V. Reis and T. Rothvoss, *The subspace flatness conjecture and faster integer programming*, in 64th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2023, November 6-9, 2023, Santa Cruz, California, USA, 2023.

[142] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.

[143] C. A. Rogers and G. C. Shephard, *The difference body of a convex body*, Archiv der Mathematik, 8 (1957), pp. 220–233.

[144] H. P. Rosenthal, *On the subspaces of L_p $(p > 2)$ spanned by sequences of independent random variables*, Israel Journal of Mathematics, 8 (1970).

[145] T. Rothvoss, *Approximating bin packing within $O(\log OPT \cdot \log\log OPT)$ bins*, in 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, IEEE, 2013, pp. 20–29.

[146] T. Rothvoss, *Constructive discrepancy minimization for convex sets*, in 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014, 2014, pp. 140–145.

[147] T. Royen, *A simple proof of the Gaussian correlation conjecture extended to multivariate gamma distributions*, arXiv: Probability, (2014).

[148] M. Rudelson, *Distances between non-symmetric convex bodies and the MM*-estimate*, Positivity, 4 (1998), pp. 161–178.

[149] M. Rudelson, *Sections of the difference body*, Discrete & Computational Geometry, 23 (1998), pp. 137–146.

[150] L. Santaló, *Un invariante afin para los cuerpos convexos del espacio de n dimensiones*, Portugaliae mathematica, 8 (1949), pp. 155–161.

[151] G. Schechtman, *More on embedding subspaces of L_p in ℓ_r^n*, Compositio Mathematica, 61 (1987), pp. 159–169.

[152] G. Schechtman, *Fourier analytic methods in convex geometry. workshop at the American Institute of Mathematics;* `http://aimpl.org/fourierconvex/1/`, 2007.

[153] R. Schneider, *Convex Bodies: The Brunn-Minkowski Theory*, Encyclopedia of Mathematics and its Applications, Cambridge University Press, 2 ed., 2013.

[154] A. Schrijver, *Theory of linear and integer programming*, Wiley-Interscience series in discrete mathematics and optimization, Wiley, 1999.

[155] Z. Šidák, *Rectangular confidence regions for the means of multivariate normal distributions*, J. Amer. Statist. Assoc., 62 (1967), pp. 626–633.

[156] J. Spencer, *Balancing games*, J. Comb. Theory, Ser. B, 23 (1977), pp. 68–74.

[157] J. Spencer, *Six standard deviations suffice*, Transactions of the American Mathematical Society, 289 (1985), pp. 679–706.

[158] J. Spencer, *Ten Lectures on the Probabilistic Method*, Society for Industrial and Applied Mathematics, 2nd edition ed., 1994.

[159] E. Steinitz, *Bedingt konvergente reihen und konvexe systeme*, Journal für die reine und angewandte Mathematik, 143 (1913), pp. 128–176.

[160] N. Stephens-Davidowitz, *On the Gaussian Measure Over Lattices*, PhD thesis, New York University, USA, 2017.

[161] S. J. Szarek and E. Werner, *A nonsymmetric correlation inequality for Gaussian measure*, Journal of Multivariate Analysis, 68 (1999), pp. 193–211.

[162] M. Talagrand, *Embedding subspaces of L_1 into ℓ_1^N*, Proceedings of the American Mathematical Society, 108 (1990), pp. 363–369.

[163] T. Tkocz, *High-dimensional phenomena: dilations, tensor products and geometry of L_1*, PhD thesis, University of Warwick, 2015.

[164] J. A. Tropp, *An introduction to matrix concentration inequalities*, Foundations and Trends in Machine Learning, 8 (2015), pp. 1–230.

[165] J. D. Vaaler, *A geometric inequality with applications to linear forms.*, Pacific Journal of Mathematics, 83 (1979), pp. 543–553.

[166] R. van Handel, *Probability in high dimension*, https://apps.dtic.mil/sti/tr/pdf/ADA623999.pdf, 2014.

[167] R. Vershynin, *High-Dimensional Probability: An Introduction with Applications in Data Science*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2018.

[168] B.-H. Vritsiou, *Regular ellipsoids and a Blaschke-Santaló-type inequality for projections of non-symmetric convex bodies*, 2023.

[169] N. Weaver, *The Kadison-Singer problem in discrepancy theory*, Discrete Mathematics, 278 (2004), pp. 227–239.

[170] A. C.-C. Yao, *Probabilistic computations: Toward a unified measure of complexity*, in Proceedings of the 18th Annual Symposium on Foundations of Computer Science, SFCS '77, USA, 1977, IEEE Computer Society, pp. 222–227.

[171] A. Zouzias, *A matrix hyperbolic cosine algorithm and applications*, in Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I, 2012, pp. 846–858.