CS5330 Research Project: Discrepancy Theory

Shawn Lim
A0184451E

May 3, 2020

1 Introduction

Given a set system (V, S), where $V = \{1, ..., n\}$ is a ground set of n elements and $S = \{S_1, ..., S_m\}, m \geq n$ is a collection of sets $S_i \subseteq V, \forall i$, discrepancy theory is concerned with finding a two-colouring $\chi : V \rightarrow \{-1, 1\}$ so that the difference between objects with different colors is small across all the sets in S. More formally, we want to find a coloring that minimizes $\max_{S_i \in S} |\Sigma_{i \in S_i} \chi(i)|$, notated disc$(S)$. A simple random coloring gives $O(\sqrt{nlogm})$ but in a well-known paper, Spencer improved this with a beautiful non-constructive proof using the pigeonhole principle showing a $O(\sqrt{n\log(m/n)})$ upper bound (Spencer, 1985).

The past decade has seen significant progress, with the first algorithms found to constructively achieve Spencer's bound. Lovett and Meka published an algorithm that involved a random walk along a polytope in $\mathbb{R}^n$ (Lovett and Meka, 2015) and soon after, Rothvoss published a stunning algorithm again achieving Spencer's bound, but is notable for its simplicity and how it further generalizes the achievement of Lovett and Meka. (Rothvoss, 2017). The achievements of these papers is what we will survey in this report.

2 Preliminaries

Before discussing the content of the algorithms, we first go over some necessary preliminaries. Let $v_1, ..., v_m \in \mathbb{R}^n$ be indicator vectors corresponding to the sets in S, then we can express the discrepancy of S as $\max_{i \in [m]} |\langle \chi, v_i \rangle|$. We write $x \sim N(0, 1)$ to denote that x is distributed according to the standard normal distribution, and $x \sim N^n(0, 1)$ to denote the n-dimensional standard normal distribution. We will use $\|.\|_2$ to denote the standard euclidean norm, and $\{e_1, ..., e_n\}$ to denote the standard basis for $\mathbb{R}^n$. In the modern context of discrepancy theory, we can actually allow $v_1, .., v_m$ to be general vectors with bounded entries (Meka 2014a); but we will assume $\|v_j\| = 1 \ \forall j$ as this normalization does not change the problem (Lovett and Meka, 2015, p.3).

2.1 Lovett and Meka

Let $V \subseteq \mathbb{R}^n$ be a subspace, and let $\{v_1, ..., v_d\}$ be an orthonormal basis for V. We write $G \sim N(V)$ to denote the standard normal distribution supported on V, so $G = G_1 v_1 + ... + G_d v_d$ where each $G_i \sim N(0, 1)$ and the G_i's are independent. This definition of $G(V)$ is invariant of our choice of basis (Lovett and Meka, 2015, p. 5). Given two normal distributions $X \sim N(\mu_1, \sigma_1^2)$ and $Y \sim N(\mu_2, \sigma_2^2)$, note that $X + Y \sim N(\mu_1 + \mu_2, \sigma_1^2, +\sigma_2^2)$.

Claim 1.1: Let $V \subseteq \mathbb{R}^n$ be a subspace and $G \sim N(V)$. Then $\forall u \in R^n, \langle G, u \rangle \sim N(0, \sigma^2)$ where $\sigma^2 \leq \|u\|_2^2$.

Proof: Let $\{v_1, ..., v_d\}$ be an orthonormal basis for V. Then $\langle G, u \rangle = (\sum_{i=1}^d \langle v_i, u \rangle \cdot G_i) \sim N(0, \sigma^2)$ where $\sigma^2 = \sum_{i=1}^d (\langle v_i, u \rangle)^2 \leq \|u\|_2^2$ □

Claim 1.2: Let $V \subseteq \mathbb{R}^n$ be a subspace, $G \sim N(V)$, and $\langle G, e_i \rangle \sim N(0, \sigma_i^2)$ (Claim 1.1). Then $\sum_{i=1}^n \sigma_i^2 = \dim(V)$.

Proof: Since $\langle G, e_i \rangle \sim N(0, \sigma_i^2)$, $E[(\langle G, e_i \rangle)^2] = \sigma_i^2$. Then $\sum_{i=1}^n \sigma_i^2 = \sum_{i=1}^n E[(\langle G, e_i \rangle)^2] = E[\sum_{i=1}^n \langle G, e_i \rangle)^2]$ $= E[\|G\|_2^2]$ (the inner product with each e_i "pulls" the ith component) $= \sum_{i=1}^d \|v_i\|_2^2 E[G_i^2]$ (orthogonality) $= \dim(V)$ (recall that each $G_i \sim N(0, 1)$). □

The following is a well known tail-bound for the standard normal distribution:

Claim 1.3: Let $X \sim N(0,1)$. Then for any $\lambda > 0$, $\Pr[|G| \geq \lambda] \leq 2\exp(-\lambda^2/2)$. (Mitzenmacher and Upfal, 2017, p.247).

We will also need the following claim, quoted directly from the paper (Lovett and Meka, 2015, p.6).

Claim 1.4: Let $X_1, ..., X_T$ be random variables. Let $Y_1, ..., Y_T$ be random variables where each Y_i is a function of X_i. Suppose that for all $1 \leq i \leq T, x_1, ..., x_{i-1} \in R$, $Y_i|(X_1 = x_1, X_2 = x_2, ..., X_{i-1} = x_{i-1})$ is Gaussian with mean zero and variance at most one (possibly different for each setting of $x_1, ..., x_{i-1}$). Then for any $\lambda > 0$, $\Pr[|Y_1 + ... + Y_T| \geq \lambda\sqrt{T}] \leq 2\exp(-\lambda^2/2)$.

2.2 Rothvoss

Here, γ_n is used to denote the measure of the n-dimensional normal distribution $N^n(0,1)$ with density $\frac{1}{(2\pi)^{n/2}}e^{\|x\|_2^2/2}$. So for a set K, $\gamma_n(K) = \Pr_{x \sim N^n(0,1)}[x \in K]$ that is to say, the measure of a set is the probability a vector distributed according to the n-dimensional standard normal distribution falls within the set (we refer to these as Gaussian vectors henceforth). We define the distance between a point and a set as $d(x, K) = \min\{y \in K : ||x - y||_2\}$ and the neighborhood of a set K as $K_\delta = \{x \in \mathbb{R}^n : d(x, K) \leq \delta\}$ (all points with distance at most δ from K, note that $K \subseteq K_\delta$).

Claim 2.1: Let $K \subseteq \mathbb{R}^n$ be a measurable set and H be a halfspace so that $\gamma_n(K) = \gamma_n(H)$. Then for any $\delta \geq 0$, $\gamma_n(K_\delta) \geq \gamma_n(H_\delta)$. (The gaussian isoperimetric inequality, Ledoux and Talagrand, 2013, p.17).

Claim 2.2: Let $\epsilon > 0$. Then for any measurable set K with $\gamma_n(K) \geq e^{-\epsilon n}$, one has $\gamma_n(K_{3\sqrt{\epsilon n}}) \geq 1 - e^{-\epsilon n}$.

Proof: Suppose we have a set K with measure $\gamma_n(K) \leq e^{-\epsilon n}$. Let $\lambda \in \mathbb{R}$ be such that the halfspace $H = \{x \in \mathbb{R}^n : x_1 \leq \lambda\}$ satisfies $\gamma_n(K) = \gamma_n(H)$. By using the one-tailed version of Claim 1.3, we have that $\Pr_{x_1 \sim N(0,1)}(x_1 \leq \frac{-3}{2}\sqrt{\epsilon n}) \leq \exp(-\frac{9}{8}\epsilon n) \leq e^{-\epsilon n}$ which implies $|\lambda| \leq \frac{3}{2}\sqrt{\epsilon n}$ and that $\gamma_n(H_{3\sqrt{\epsilon n}}) \leq 1 - e^{-\epsilon n}$. This implies that if K has measure $\gamma_n(K) \geq e^{-\epsilon n}$, the corresponding halfspace satisfies $\gamma_n(H_{3\sqrt{\epsilon n}}) \geq 1 - e^{-\epsilon n}$ (by symmetry). By Claim 2.1, we then have that $\gamma_n(K_{3\sqrt{\epsilon n}}) \geq \gamma_n(H_{3\sqrt{\epsilon n}}) \geq 1 - e^{-\epsilon n}$ $\square$.

Intuitively, the point of Claim 2.1 and 2.2 is that for any set whose gaussian measure is not too small, a gaussian vector will be near the set with good probability. We call a set K convex if $\forall x, y \in K, \lambda \in [0,1], \lambda x + (1-\lambda)y \in K$. We call a set K symmetric if $x \in K \leftrightarrow -x \in K$. Let us call a set of the form $\{x \in \mathbb{R}^n : |\langle v, x\rangle| \leq \lambda\}$ a strip. The following claim, from a lemma of Sidak (Sidak, 1967) and Khatri (Kahtri, 1967) says that the intersection of a set with a strip, does not cause the measure to decrease by much:

Claim 2.3: Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body and $S \subseteq \mathbb{R}^n$ be a strip. Then $\gamma_n(K \cap S) \geq \gamma_n(K) \cdot \gamma_n(S)$

In preparation for a union bound argument later, for some $\epsilon \in [0, \frac{1}{2}]$, we want to bound the number of subsets of n elements of size ϵn. we present the following combinatorial argument (Mitzenmacher and Upfal, p.272) with relevant additions.

Claim 2.4: The number of subsets of n elements of size $\epsilon n \leq (e^{n\frac{3}{2}\epsilon \log_2(\frac{1}{\epsilon})})$.

Proof: $\binom{n}{\epsilon n}\ \epsilon^{\epsilon n}(1-\epsilon)^{(1-\epsilon)n} \leq \sum_{k=0}^{n} \epsilon^k(1-\epsilon)^{n-k} \leq (\epsilon + (1-\epsilon))^n = 1$ where the last inequality follows from the binomial theorem. Noting that $\log_2(\frac{1}{\epsilon}) > \log_2(\frac{1}{1-\epsilon})$ for $\epsilon \in [0, \frac{1}{2}]$, this implies that:

$$\binom{n}{\epsilon n} \leq \epsilon^{-\epsilon n}(1-\epsilon)^{-(1-\epsilon)n} \leq 2^{\epsilon n \log_2(\frac{1}{\epsilon})} \cdot 2^{(1-\epsilon)n \log_2(\frac{1}{1-\epsilon})} \leq 2^{n(\epsilon \log_2(\frac{1}{\epsilon}) + (1-\epsilon)\log_2(\frac{1}{1-\epsilon}))} \leq e^{n(\frac{3}{2}\epsilon \log_2(\frac{1}{\epsilon}))} \qquad \square$$

Lastly, we need the following result concerning optimal solutions to convex functions in convex sets that we quote directly from the paper (Rothvoss, 2017, p.5). Note that if a function g is strictly convex, g satisfies $g(\lambda x + (1-\lambda)y) < \lambda g(x) + (1-\lambda)g(y)$ for x, y in g's domain.

Claim 2.5: Let $P, Q \subseteq \mathbb{R}^n$ be convex sets and let $g : \mathbb{R}^n \to \mathbb{R}$ be a strictly convex function. If x^* optimal for $\min\{g(x) : x \in P \cap Q\}$ and x^* lies in the interior of Q, then x^* is also optimal for $\min\{g(x) : x \in P\}$.

3 Main theorems

In both algorithms the goal is not a full coloring $\epsilon \in \{-1,1\}^n$, but instead a partial coloring $\epsilon \in \{-1,0,1\}^n$ where (a) there is a lower-bound on the number of non-zero entries and (b) the partial coloring has low discrepancy. We can induce on the remaining colors and if we find a constant fraction of n colors at each step, then in $O(logn)$ iterations we have a full coloring. Both algorithms achieve this in polynomial time.

3.1 Constructive discrepancy by walking on edges

We first survey the algorithm of Lovett and Meka. Given the set system (V, S), we can rephrase our goal of low discrepancy in a geometric way, where instead of finding a partial coloring ϵ so that $|\langle v_j, \epsilon\rangle| = O(\sqrt{n})\ \forall j$ (v_j are the indicator vectors for our set system), we look for lattice points with sufficiently many non-zero entries inside a symmetric convex set $P = \{x \in \mathbb{R}^n : |\langle v_j, x\rangle| \leq \Delta,\ \forall j \in [m],\ |x_i| \leq 1\ \forall i \in [n]\}$ where $\Delta = O(\sqrt{n})$ (Meka 2014c). The description of the set has two parts; the first ensuring the points inside satisfy our low discrepancy requirement, while the second ensures the colors are at least in the range [-1, 1]. Specifically, what we need is a point inside this set with $\Omega(n)$ non-zero integral coordinates.

Instead of actually finding such a point, we search for a point far from the origin but still inside the set; the idea being that such a point will have many coordinates close to $\{-1,1\}$, and we do this via a random walk, taking small gaussian steps at each iteration. To ensure we remain inside the polytope P, we make use of a parameter δ; once we become close to either a variable constraint ($|x_i| \geq 1-\delta$) or a discrepancy constraint ($|\langle v_j, x\rangle| \geq \Delta - \delta|$), we constrain future steps inside the edge of the polytope we hit (hence the name walking on "edges"). Here is the formal statement of their algorithm:

3.1.1 The algorithm

Let $(X_0 = x_0, ..., X_T = x)$ be a series of T steps where X_0 is our starting point and X_T is the final result returned. γ is a parameter related to the size of the step we take; set $\gamma > 0$, such that $\delta = O(\gamma\sqrt{log(nm/\gamma)})$ (we will assume $\delta < 0.1$) and let T = K_1/γ^2 where $K_1 = 16/3$. Let $c_1, ..., c_m$ be the discrepancy constraints for each set. Now, at an intermediate step t, let:

- $C_t^{\text{var}} = \{i \in [n] : |(X_{t-1})_i| \geq 1-\delta\}$ be the set of "nearly hit" variables,
- $C_t^{\text{disc}} = \{j \in [m] : |\langle X_{t-1}, v_j\rangle| \geq c_j - \delta\}$ be the set of "nearly hit" constraints and
- $V_t = \{u \in \mathbb{R}^n : u_i = 0\ \forall i \in C_t^{\text{var}}\ \langle u, v_j\rangle = 0\ \forall j \in C_t^{\text{disc}}\}$ be the subspace orthogonal to both.

Then at each step, simply set $X_t = X_{t-1} + \gamma U_t$ where $U_t \sim N(V_t)$. We have the following guarantee:

Theorem 1: Let $\{v_1, ..., v_m\} \in \mathbb{R}^n$ be vectors, and $x_0 \in \{-1,1\}^n$ be a "starting point". Let $c_1, ..., c_m \geq 0$ be thresholds such that $\sum_{j=1}^m \exp(-c_j^2/16) \leq n/16$. Let $\delta > 0$ be a small approximation parameter. Then there exists an efficient randomized algorithm which with probability at least 0.1 finds a point $x \in [-1,1]^n$ such that:
(i) $|x_i| \geq 1-\delta$ for at least $n/2$ indices $i \in [n]$ (there are enough new colors assigned).
(ii) $|\langle x - x_0, v_j\rangle| \leq c_j\|v_j\|_2$ (the discrepancy does not increase by too much).

To prove Theorem 1, our main goals are to show that with probabiliy at least 0.1, (1): all steps remain inside the polytope and that (2): $(X_T)_i \geq 1-\delta$ for at least $n/2$ indices.

3.1.2 All steps remain inside the polytope

Once a variable or discrepancy constraint becomes "nearly hit" at some time t, drawing the direction we walk in from the orthogonal subspace ensures that the constraint remains unchanged thereafter. So what remains is to bound the probability that a constraint is violated at the same time it becomes "nearly hit".

Suppose a variable constraint i is violated at time t: since the constraint is not tight at time $(t-1)$, $|(X_{t-1})_i| \leq 1-\delta$, but for the constraint to be violated we also have that $|(X_t)_i| \geq 1$, which implies that $\gamma|(U_t)_i| \geq \delta$. Analogously for a discrepancy constraint j, we have that $\gamma|\langle U_t, v_j\rangle| \geq \delta \implies |\langle U_t, v_j\rangle| \geq \delta/\gamma$. We can similarly express the violation of the variable constraint as an inner product: $|\langle U_t, e_i\rangle| \geq \delta/\gamma$.

Let $W = \{e_1, ..., e_n, v_1, ..., v_m\}$. Then the probability at least one constraint is violated at time t is $\sum_{w \in W} \Pr(|\langle U_t, w\rangle| \geq \delta/\gamma)$. By Claim 1.3, we have that $\Pr(|\langle U_t, w\rangle| \geq \delta/\gamma) \leq 2\exp(-(\delta/\gamma)^2)/2)$. Since we set $\delta = O(\gamma\sqrt{log(nm/\gamma)})$, we can write $\delta/\gamma = \sqrt{2Clog(nm/\gamma)}$ for some constant C. Then using a union bound over all times t and all $n+m$ vectors in W (bounded above by nm), we have that:
$\Pr(\text{any constraint is violated at any time t}) \leq T \cdot (nm) \cdot 2\exp(-(\delta/\gamma)^2)/2) \leq T \cdot (nm) \cdot \frac{2\gamma^C}{(nm)^C}$

Noting that $T = \frac{16}{3\gamma^2}$ and $\gamma < 1$, when C is suff. large: $T \cdot (nm) \cdot \frac{2\gamma^C}{(nm)^C} = \frac{\frac{32}{3}\gamma^{C-2}}{(nm)} \cdot \frac{1}{(nm)^{C-2}} \leq \frac{1}{(nm)^{C-2}}$

So we conclude that all steps remain inside the polytope with high probability $\square$

3.1.3 Sufficiently many coordinates are close to $\{-1, 1\}$

The outline for this part of the argument is as follows: we can show that $|C_t^{\text{disc}}|$ is generally small. Thus if $\dim(V_t)$ is small, this implies $|C_t^{\text{var}}|$ is big, which is what we want. Otherwise, if $|C_t^{\text{var}}|$ is small, then $\dim(V_t)$ is large and we argue this implies $\|X_t\|_2^2$ is likely to increase significantly. If X_t remains inside the polytope as its norm increases, together this implies that more variable constraints are becoming tight and thus $|C_t^{\text{var}}|$ cannot remain small for many steps, which completes our proof for Theorem 1. We first show that $|C_t^{\text{disc}}|$ is small $\forall t$.

Claim 1.5: $E[|C_T^{\text{disc}}|] < n/4$.
Proof: The idea is that we separate constraints into two groups; constraints that easily become tight, and constraints that do not. Let $J = \{j : c_j \leq 10\delta\}$. Since $\sum_{j=1}^m \exp(-c_j^2/16) \leq n/16$ and $\delta < 0.1$, we have that: $n/16 \geq \sum_{j \in J} \exp(-c_j^2/16) \geq |J|\exp(-100\delta^2/16) \geq |J|e^{-1/16} > |J| \cdot 9/10 \implies |J| < (10/9)n/16 < 1.2n/16$.

Now for $j \notin J$: if $j \in C_T^{\text{disc}}$, then $|\langle X_T - x_0, v_j\rangle| \geq c_j - \delta > 0.9c_j$. From the algorithm, $X_T = x_0 + \gamma(U_1, ..., U_T)$, thus the previous inequality is equivalent to $|\langle \gamma U_1, v_j\rangle| + ... + |\langle \gamma U_T, v_j\rangle| > 0.9c_j$. Let $Y_t = \langle U_t, v_j\rangle$. From Claim 1.1, $Y_t \sim N(0, \sigma^2), \sigma^2 \leq 1$. We can then rewrite the inequality again as $|Y_1 + ... + Y_T| \geq 0.9c_j/\gamma$. $U_1, ..., U_T$ are a sequence of random variables, and each $Y_t|U_1, ..., U_{t-1}$ is Gaussian with mean zero and variance at most one. Thus by Claim 1.4 (noting again that $T = K_1/\gamma^2, K_1 = 16/3$):

$Pr(|Y_1 + ... + Y_T| \geq \frac{0.9c_j}{\gamma \cdot \sqrt{T}} \cdot \sqrt{T}) \leq 2\exp(-(0.9c_j)^2/2\gamma^2 T) = 2\exp(-(0.9c_j)^2/2K) < 2\exp(-c_j^2/16)$.
By assuming that all constraints in J are hit, and the fact that $\sum_{j=1}^m \exp(-c_j^2/16) \leq n/16$ it follows that:

$E[|C_T^{disc}|] \leq |J| + \sum_{j \notin J} \Pr(j \in C_T^{disc}) \leq |J| + 2\sum_{j=1}^m \exp(-c_j^2/16) \leq 1.2n/16 + 2n/16 < n/4$ $\square$

Claim 1.6: $E[||X_T||_2^2] \leq n$.
Proof: We do this by showing $E[(X_T)_i^2] \leq 1\ \forall i$. Conditioning on the first t such that variable i becomes tight: $E[(X_T)_i^2] = \Pr(i \notin C_T^{var})E[(X_T)_i^2|i \notin C_T^{var}] + \sum_{t=1}^T Pr(i \in C_t^{var} \backslash C_{t-1}^{var})E[(X_t)_i^2|i \in C_t^{var} \backslash C_{t-1}^{var}]$

The support for i in cases where the constraint is not tight is $\leq 1 - \delta$ so clearly $E[(X_T)_i^2|i \notin C_T^{var}] \leq 1$. On the other hand, we can upper-bound $(X_t)_i = (X_{t-1}) + \gamma(U_t)_i \leq 1 - \delta + \gamma(U_t)_i$ when X_{t-1} is not tight. Then

$E[(X_t)_i^2|i \in C_t^{var} \backslash C_{t-1}^{var}] \leq E[((1-\delta) + \gamma(U_t)_i)^2] \leq (1-\delta)^2 + 2\gamma(1-\delta)E[(U_t)_i] + \gamma^2 E[(U_t)_i^2] \leq 1 - \delta + \gamma \leq 1$

Since $(U_t)_i = \langle U_t, e_i\rangle \sim N(0, \sigma^2), \sigma^2 \leq 1$ by Claim 1.1 $\square$

Claim 1.7: $E[C_T^{var}] \geq 0.56n$.
Proof: First we compute the average norm of X_t.
$E[\|X_t\|_2^2] = E[\|X_{t-1} + \gamma U_t\|_2^2] = E[\|X_{t-1}\|] + \gamma^2 E[\|U_{t-1}\|_2^2) + 2\gamma E[\langle X_{t-1}, U_t\rangle] = E[\|X_{t-1}\|_2^2] + \gamma^2 E[\dim(V_t)]$

Here, use the fact that $E[\|U_{t-1}\|_2^2] = E[\sum_{i=1}^n \langle U_{t-1}, e_i\rangle^2] = \sum_{i=1}^n E[\langle U_{t-1}, e_i\rangle^2] = \dim(V_t)$ by Claim 1.2 and $E[\langle X_{t-1}, U_t\rangle] = E[E[\langle X_{t-1}, U_t\rangle|X_{t-1}]] = \sum_{x_{t-1}} \sum_i \Pr(X_{t-1} = x_{t-1})(x_{t-1})_i \cdot E[(U_t)_i|X_{t-1} = x_{t-1}] = 0$.

Then applying Claim 1.5, we have that:

$n \geq E[\|X_T\|_2^2] \geq \gamma^2 \sum_{t=1}^T E[\dim(V_t)] \geq \gamma^2 T \cdot E[\dim(V_T)] = K_1 E[\dim(V_T)] = K_1 E[(n - |C_T^{var}| - |C_T^{disc}|)]$

Then: $n \geq K_1 n - E[|C_T^{var}|] - E[|C_T^{disc}|]$, thus $E[|C_T^{var}|] \geq n(1 - 3/16) - n/4$ (Claim 1.6) $> 0.56n$ $\square$

3.1.4 Conclusion

We finish with an argument using Markov bounds. Suppose $\Pr[|C_T^{var} \geq 0.5n] = \alpha$, then: $E[|C_T^{var}|] \leq \alpha \cdot n + (1-\alpha) \cdot 0.5n = 0.5n + 0.5\alpha n$. $\alpha < 0.12$ would imply $E[|C_T^{var}|] < 0.5n + 0.06n = 0.56n$, contradicting Claim 1.7. Thus with probability at least 0.12, $|C_T^{var}| \geq 0.56n$. Coupled with the fact that all steps are inside the polytope with probability at least $1/(nm)^{C-2}$, the algorithm succeeds with probability at least $0.12 - 1/(nm)^{C-2} > 0.1$ $\square$

For completeness, we sketch some of the finer technical details we could not cover in depth. We run the algorithm 2log$n$ times, inducing on the remaining "non-tight" variables at each iteration and appropriately concatenating partial colorings we find. Choosing $8\sqrt{\log(m/n)}$ for the thresholds at each iteration will meet the conditions and satisfy a $O(\sqrt{n\log(m/n)})$ discrepancy (note that n is decreasing across iterations!). To improve accuracy, simply run the algorithm multiple times for each set of partial colorings (they recommend $O(\log n)$ iterations in the paper), and randomized rounding is used to find the final full coloring (Lovett and Meka, 2015, p.9).

3.2 Constructive discrepancy minimization for convex sets

3.2.1 The algorithm

Rothvoss takes a similar geometric approach, but his method for finding a good partial coloring is stunningly simple. Let K be a symmetric convex body. Now, simply generate a random Gaussian vector $x^* \sim N^n(0,1)$, then compute and return $y* = \text{argmin}\{\|x^* - y\|_2 \; |y \in K \cap [-1,1]^n\}$. He provides the following guarantee:

Theorem 2: Let $0 < \epsilon \leq \frac{1}{9000}$ be a constant and let $\delta = \frac{3}{2}\epsilon\log_2(\frac{1}{\epsilon})$. Suppose that $K \subseteq \mathbb{R}^n$ is a symmetric convex body with $\gamma_n(K) \geq e^{-\delta n}$. Choose a random Gaussian vector $x^* \sim N^n(0,1)$ and let y^* be the point in $K \cap [-1,1]^n$ be the point that minimizes $\|x^* - y^*\|_2$. Then with probability $1 - e^{-\Omega(n)}$, y^* has at least ϵn many coordinates i with $y_i^* \in \{-1, 1\}$.

3.2.2 Proof of main theorem

The outline of the argument is as follows: we can show that x^* is at least $\Omega(\sqrt{n})$ away from the hypercube $[-1,1]^n$. On the other hand, we can define a set, whose size is related to the number of integral coordinates of our solution y^*, and has measure inversely proportional to its size, i.e. when this set is large, many coordinates of y^* are integral, but the set has small measure. The distance of x^* to this set is related to the distance of x^* from $K \cap [-1,1]^n$. Thus if this set is too small, its measure is large, and then by the Gaussian isoperimetric inequality, x^* will be too near to this set, and by extension be too near to $K \cap [-1,1]^n$. This implies the set is at least of a certain size, ensuring y^* has $\Omega(n)$ integral coordinates.

We begin with the argument showing x^* is at least $\Omega(\sqrt{n})$ away from $[-1,1]^n$. It can be directly computed that for a component of x^*, $\Pr_{x_i \sim N(0,1)}(|x_i| \geq 2) > 1/25$. Treating each coordinate as an indicator, we can bound the number of such coordinates about $n/25$ with a standard chernoff bound and show that $d(x^*, [-1,1]^n) \geq \sqrt{\frac{n}{25} \cdot (2-1)^2} = \frac{1}{5} \cdot \sqrt{n}$ with probability $1 - e^{-\Omega(n)}$.

Now we define the set mentioned above. $K \cap [-1,1]^n$ itself has only a small Gaussian measure, but we can consider a super-set containing it: $K(I^*) := K \cap \{x \in \mathbb{R}^n : |x_i| \leq 1 \; \forall i \in I^*\}$ where $I^* := I^*(x^*) := \{i \in [n] \; |y_i^* \in \{-1,1\}\}$. Crucially, we can argue that $d(x^*, K \cap [-1,1]^n) = d(x^*, K(I^*))$, by defining $P := K(I^*), Q := \{x \in \mathbb{R}^n : |x_i| \leq 1 \; \forall i \notin I^*\}$; then $K \cap [-1,1]^n = P \cap Q$ and since $\|x^* - y\|_2$ is a strictly convex function, we can apply Claim 2.5.

Notice that the size of the set I^* is equal to the number of integral coordinates of our result y^*. Our guarantee was ϵn many integral coordinates, thus we want to lower-bound the measure of the set $K(I^*)$ when $|I^*| \leq \epsilon n$. $K(I^*)$ is the intersection of K with strips, and we have the following estimate for the measure of a strip: $\gamma_n(\{x \in \mathbb{R}^n : |x_i| \leq \lambda\}) \geq 1 - e^{-\lambda^2/2}$ (Rothvoss, p.4) . Now noting that our theorem assumed $\gamma_n(K) \geq e^{-\delta n}$ and our choice of δ ensures $\epsilon \leq \delta$, by Claim 2.4:

$$\gamma_n(K(I^*)) \geq \gamma_n(K) \cdot \prod_{i \in I^*} \gamma_n(\{x \in \mathbb{R}^n : |x_i| \leq 1\}) \geq \gamma_n(K) \cdot e^{-|I^*|/2} \geq e^{-\delta n} \cdot e^{-(\epsilon n)/2} \geq e^{-2\delta n}$$

Now by Claim 2.2, $\gamma_n(K(I^*)) \geq e^{-2\delta n}$ implies $\gamma_n(K(I^*)_{3\sqrt{2\delta n}}) \geq 1 - e^{-2\delta n}$. Thus with high probability, a random Gaussian vector is in a $3\sqrt{2\delta n} < \frac{1}{5}\sqrt{n}$ neighborhood of $K(I^*)$. However, I^* as defined above was dependent on x^*, so now we instead consider all possible subsets of $[n]$ of size $\leq \epsilon n$. Let $B = \cap_{|I| \leq \epsilon n}(K(I)_{3\sqrt{2\delta n}})$, so B is the set of points common to all $3\sqrt{2\delta n}$ neighborhoods for any set $K(I)$ such that $|I| \le \epsilon n$. By our choice of δ and Claim 2.4, there are at most $e^{\delta n}$ many subsets I such that $|I| \le \epsilon n$. Finally with a union-bound:

$$\gamma_n(B) = 1 - \gamma_n(\cup_{|I|\le\epsilon n}\mathbb{R}^n \backslash K(I)_{3\sqrt{2\delta n}}) \ge 1 - \sum_{|I|\le\epsilon n}\gamma_n(\mathbb{R}^n\backslash K(I)_{3\sqrt{2\delta n}}) \ge 1 - e^{-\delta n}e^{-2\delta n} \ge 1 - e^{-\delta n} \ \square$$

Thus a random Gaussian vector is within $3\sqrt{2\delta n}$ from all sets $K(I), |I| \le \epsilon n$ with probability at least $1 - e^{-\delta n}$, while being $\frac{1}{5}\sqrt{n} > 3\sqrt{2\delta n}$ away from the hypercube with probability at least $1 - e^{-\Omega n}$. This implies that with high probability, $|I^*| > \epsilon n$ $\square$

3.2.3 Conclusion

At this point, we would like to compare this result with the achievement of Lovett and Meka, in addition to providing an outline for how Theorem 2 can be used to find a good coloring. The main requirement for Theorem 1 is that the thresholds we specify $c_1, ..., c_m \ge 0$ satisfy $\sum_{j=1}^m \exp(-c_j^2/16) \le \frac{n}{16}$. However this condition does not appear to apply to arbitrary convex sets, and in particular, convex sets with large measure may still be unable to satisfy this condition (Rothvoss, p.2).

On the other hand, Rothvoss's result applies to **any** symmetric convex body with sufficiently large measure, and he shows in his paper (Rothvoss p.7) that a polytope meeting the requirements of Theorem 1 meets the conditions for Theorem 2. Thus we can view his work as further generalizing the contexts in which we can minimize discrepancy, in some sense containing Lovett and Meka's result.

To apply Theorem 2 to find a good coloring, we generate the random Gaussian vector $x*$ and compute y_1 to find the first partial coloring. Subsequently for $t \ge 2$, let $U = \{x \in \mathbb{R}^n : x_i = 0 \text{ if } (y_{t-1})_i \in \{-1, 1\}\}$. Set the previous y_{t-1} as the new center and compute $y_t \in (y_{t-1} + K \cap U)$ where $(y_{t-1} + K \cap U)$ is the set $\{x \in K \cap U : y_{t-1} + x\}$. Intuitively, Theorem 2 can be extended to allow centers that are not the origin, and it can be shown that the set $K \cap U$ has sufficiently large measure (Rothvoss p.8).

Taking the intersection with the subspace U is essentially inducing on the components of y_{t-1} that are not integral and since a constant fraction of components become integral at each step, the algorithm finds a full coloring in $O(logn)$ iterations.

4 The Beck-Fiala Conjecture

To conclude our survey, we consider two of the longstanding open problems in the field, concerning the special case where each element in our ground set V appears at most t times. The first result in the field is the well known Beck-Fiala theorem where it was proven that for a set system (V, S) satisfying this condition, we have $\text{disc}(S) < 2t - 1$ (Beck and Fiala, 1981). A beautiful proof of this fact using simple linear algebra can be found in the book of Chazelle (Chazelle, 2001, p.10). Beck and Fiala conjectured that the correct bound was actually $O(\sqrt{t})$ but the best known bound remains at $O(\sqrt{tlogn})$ (Banaszczyk, 1998) given by Banaszczyk using arguments from convex geometry. This remains a significant open problem in discrepancy theory today.

Interestingly, finding an algorithm that constructively achieves Banaszczyk bounds is also currently an open problem that is still being actively looked at (Dadush, Garg, Lovett and Nikolov, 2016). Here we will briefly look at how Lovett and Meka's algorithm can be applied to achieve the best known constructive bound of $O(\sqrt{t}\log n)$. In the general case, they suggest the thresholds $8\sqrt{\log(m/n)}$ for all the c_i, but in this particular setting, they exploit the "sparseness" of the sets to set lower thresholds.

If each element appears at most t times, it is straightforward to see that for our identity vectors $v_1, ..., v_m$, $\sum_{j=1}^m \|v_j\|_2^2 \le nt$. This implies that you have at most $n/2^r$ vectors such that $\|v_j\|_2^2 \in [2^r t, 2^{r+1}t]$. Then by picking the thresholds $c_j = C\sqrt{t}/\|v_j\|_2$, and considering these intervals of size 2^r, we have that:

$$\sum_{j=1}^m \exp(-c_j^2/16) = \sum_{j=1}^m \exp(-C^2 t/16\|v_j\|_2) < \sum_{r=0}^\infty \frac{n \cdot \exp(-C^2/(16 \cdot 2^{r+1}))}{2^r} < n/16$$

When C is large enough. Now each time we apply the algorithm to find a partial coloring, we have that $|\langle x_t - x_{t-1}, v_j\rangle| < c_j\|v_j\|_2 = C\sqrt{t}$. Since we expect $> n/2$ coordinates to be filled each iteration, the algorithm takes $O(\log n)$ iterations and thus the total discrepancy does not exceed $O(\sqrt{t}\log n)$ $\square$

5 Bibliography

1. Spencer, J. (1985). Six standard deviations suffice. Transactions of the American mathematical society, 289(2), 679-706.

2. Lovett, S., & Meka, R. (2015). Constructive discrepancy minimization by walking on the edges. SIAM Journal on Computing, 44(5), 1573-1582.

3. Rothvoss, T. (2017). Constructive discrepancy minimization for convex sets. SIAM Journal on Computing, 46(1), 224-234.

4. Bansal, N. (2010, October). Constructive algorithms for discrepancy minimization. In 2010 IEEE 51st Annual Symposium on Foundations of Computer Science (pp. 3-10). IEEE.

5. Beck, J., & Fiala, T. (1981). "Integer-making" theorems. Discrete Applied Mathematics, 3(1), 1-8.

6. Banaszczyk, W. (1998). Balancing vectors and Gaussian measures of n-dimensional convex bodies. Random Structures & Algorithms, 12(4), 351-360.

7. Dadush, D., Garg, S., Lovett, S., & Nikolov, A. (2016). Towards a constructive version of Banaszczyk's vector balancing theorem. arXiv preprint arXiv:1612.04304.

8. Khatri, C. G. (1967). On certain inequalities for normal distributions and their applications to simultaneous confidence bounds. The Annals of Mathematical Statistics, 1853-1867.

9. Šidák, Z. (1967). Rectangular confidence regions for the means of multivariate normal distributions. Journal of the American Statistical Association, 62(318), 626-633.

10. Chazelle, B. (2001). The discrepancy method: randomness and complexity. Cambridge University Press.

11. Mitzenmacher, M., & Upfal, E. (2017). Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis. Cambridge university press.

12. Ledoux, M., & Talagrand, M. (2013). Probability in Banach Spaces: isoperimetry and processes. Springer Science & Business Media.

13. Meka, Raghu. (2014). Discrepancy and beating the Union Bound. Retrieved April 28, 2020, from https://windowsontheory.org/2014/02/07/discrepancy-and-beating-the-union-bound/

14. Meka, Raghu. (2014). Discrepancy Bounds from Convex Geometry. Retrieved April 28, 2020, from https://windowsontheory.org/2014/02/17/discrepancy-bounds-from-convex-geometry/

15. Meka, Raghu. (2014). Discrepancy and Rounding Linear Programs. Retrieved April 28, 2020, from https://windowsontheory.org/2014/02/27/discrepancy-and-rounding-linear-programs/