Hereditary discrepancy and factorization norms

organized by
Aleksandar Nikolov and Kunal Talwar

Workshop Summary

Background and Goals

This workshop was devoted to the application of methods from functional analysis and asymptotic convex geometry to combinatorial discrepancy theory. The central quantity of interest in combinatorial discrepancy theory is the *discrepancy* of a system of subsets $\mathcal{S} = \{S_1, \ldots, S_m\}$ of the ground set $[n]$, defined as

$$\mathrm{disc}(\mathcal{S}) \stackrel{def}{=} \min_{x \in \{-1,1\}^n} \max_{i=1}^{m} \Big|\sum_{j \in S_i} x_j\Big|.$$

In this context, a vector $x \in \{-1,1\}^n$ is called a coloring. The definition of discrepancy extends in a natural way to matrices: we define the discrepancy of an $m \times n$ real matrix A as $\min_{x \in \{-1,1\}^n} \|Ax\|_\infty$. These two definitions coincide when we take A to be the incidence matrix of $\mathcal{S}$. Matrix discrepancy is also closely related to vector balancing problems, in which we are given two Banach spaces X, Y defined on $\mathbb{R}^m$, and we are interested in the vector balancing constant

$$\alpha(X,Y) = \sup\left\{ \min_{\varepsilon_1,\ldots,\varepsilon_n} \Big\|\sum_{i=1}^{n} \varepsilon_i v_i\Big\|_Y : n \in \mathbb{N}, v_1, \ldots, v_n \in \mathbb{R}^m, \max_{i=1}^{n} \|v_i\|_X \le 1 \right\}.$$

Hereditary discrepancy is a robust variant of discrepancy, introduced by Lovász, Spencer, and Vesztergombi [LSV]. For a set system, it is defined as the maximum attainable discrepancy over all restricted subsystems obtained by deleting elements of $[n]$. For a set $J \subseteq [n]$ and an $m \times n$ real matrix A, let A_J represent the submatrix consisting of the columns of A indexed by J. Then, for a matrix A we define

$$\mathrm{herdisc}(A) \stackrel{def}{=} \max_{J \subseteq [n]} \mathrm{disc}(A_J).$$

For a set system $\mathcal{S}$, $\mathrm{herdisc}(\mathcal{S}) \stackrel{def}{=} \mathrm{herdisc}(A)$, where A is the incidence matrix of $\mathcal{S}$.

Hereditary discrepancy has deep connections to uniformity of distributions and a number of applications to theoretical computer science. However, while it is generally more tractable than discrepancy, estimating it is still challenging. Recent work by Nikolov and Talwar showed that for any $m \times n$ matrix A

$$\mathrm{herdisc}(A) \lesssim \sqrt{\log m} \cdot \gamma_2(A); \qquad (1)$$

$$\gamma_2(A) \lesssim \log(\mathrm{rank}\ A) \cdot \mathrm{herdisc}(A). \qquad (2)$$

Here $\gamma_2(A) = \inf\{\|B\|\|C\| : A = BC\}$, where we treat A as an operator from ℓ_1^n to ℓ_∞^m, and we treat B as an operator from ℓ_1 to ℓ_2, and C as an operator from ℓ_2 to ℓ_∞. This is a classical factorization norm which measures the "cost" of factoring A through Hilbert space. See [TJ-book] for extensive background on such norms. The inequalities (1) and (2) allow translating estimates on the γ_2 norm directly into estimates on hereditary discrepancy. This is a powerful connection, since γ_2 satisfies many nice properties, which often makes it much easier to estimate than discrepancy. This technique has led to easier proofs of classical results, and a much better understanding of the discrepancy of some natural collections of sets, most prominently sets induced by axis-aligned boxes in high dimension and by homogeneous arithmetic progressions.

The goal of this workshop was to refine the connections between functional analysis and convex geometry, on one hand, and combinatorial discrepancy, on the other, to explore further the implications of such techniques to both fields, and to generalize and tighten the results alluded to above.

Structure of the Workshop and Summary of Talks

The structure of the workshop was to have 1-3 morning talks and to break out into working groups in the afternoon. Each working group consisted of 3-4 people and focused on a particular problem. Oral reports of the results from each working group were given the next morning, before talks started. The afternoon of the first day of the workshop was an exception, and was devoted to identifying a list of open problems in the field. This open problem session was attended by the all workshop participants.

The morning talks given during the workshop were as follows:

- (Monday) Aleksandar Nikolov gave some background on combinatorial discrepancy, and sketched the proofs of the inequalities (1) and (2); Nicole Tomczak-Jaegermann gave background on operator ideal norms, and the γ_2 norm in particular;
- (Tuesday) Nikhil Bansal talked about discrepancy minimization algorithms, and algorithmic proofs of upper bounds in combinatorial discrepancy; Daniel Dadush talked about Banaszczyk's (non-constructive) proof of his vector balancing result [bana], which is used in the proof of (1);
- (Wednesday) Dmitriy Bilyk gave an overview of measure-theoretic discrepancy and its connections to combinatorial discrepancy, numerical integration, and compressed sensing; Daniel Dadush talked about the equivalence of Banaszczyk's vector balancing result mentioned above, and the existence of a subgaussian distribution supported on the set $\{\pm v_1 \pm \ldots \pm v_n\}$ for vectors $v_1, \ldots, v_n$ with small ℓ_2 norm;
- (Thursday) Thomas Rothvoss presented an algorithmic proof of Giannopoulos's vector balancing result [giann]; Mohit Singh talked about applications of the Lovett-Meka discrepancy minimization algorithm to sampling from a negatively associated distribution on the vertices of a matroid polytope; Esther Ezra talked about bounds on the discrepancy of random set systems with bounded degree.
- (Friday) Shachar Lovett talked about connections between combinatorial discrepancy and the log-rank conjecture in communication complexity;

Working Groups

An open problem session was organized in the afternoon of the first day of the workshop. A list of 14 open problems related to combinatorial discrepancy and vector balancing questions was compiled, and is available at the AIM website. Below we summarize the discussion and results from the working groups on these and related problems.

Vector Balancing.

Several working group discussions were related to vector balancing problems, and specifically to characterizing the vector balancing constants $\alpha(X,Y)$ in terms of geometric properties of X and Y, similarly to the characterization of hereditary discrepancy in terms of the γ_2 norm. The first discussion of the working group was dedicated to presenting and strengthening a lower bound of Banaszczyk [Bana93]. In the strengthened form, this result states that for any Banach spaces X and Y on $\mathbb{R}^m$, with unit balls respectively B_X and B_Y,

$$\alpha(X,Y) \gtrsim \max_{k=1}^{m} \max_{W:\dim(W)=k} \min_{E:B_X\cap W\subseteq E} \sqrt{k}\frac{\mathrm{vol}_W(E)^{1/k}}{\mathrm{vol}_W(B_Y\cap W)^{1/k}}, \tag{3}$$

where W ranges over linear subspaces of $\mathbb{R}^m$, E ranges over ellipsoids inside W, and the $\mathrm{vol}_W(\cdot)$ is the Lebesgue measure inside W. The group discussed an approach to proving that the inequality (3) can be reversed, up to factors polylogarithmic in m. The approach is based on giving lower bounds on the Gaussian measure of B_Y in terms of the right hand side of (3), and then using Giannopoulos's vector balancing result [giann]. The following conjecture was shown to imply that (3) is tight:

Conjecture 1. *Let X be a Banach space on $\mathbb{R}^m$ with unit ball B_X, and define*

$$S(X) \stackrel{def}{=} \max_{S\subseteq[m]} \sqrt{|S|}\mathrm{vol}_{W(S)}(B_X\cap W(S)),$$

where $W(S) = \mathrm{span}(e_i : i \in S))$ and e_i is the i-th standard basis vector of $\mathbb{R}^m$. Then $\mathbb{E}\|G\|_X = (\log m)^{O(1)}S(X)$, where G is a standard Gaussian in $\mathbb{R}^m$.

If true, Conjecture 1 would be an analogue of the Milman-Pisier volume number theorem [MP-volnum] for coordinate subspaces.

Moreover, the tightness of (3) was shown when X is an m-dimensional Hilbert space. It's easy to see that the tightness of (3) for all m-dimensional Hilbert spaces is equivalent to the tightness for $X = \ell_2^m$. Then it is a consequence of the volume number theorem that that for any Banach space Y on $\mathbb{R}^m$

$$\mathbb{E}\|G\|_Y \lesssim (\log m)^2 \max_{k=1}^{m} \max_{W:\dim(W)=k} \min_{E:B_X\cap W\subseteq E} \frac{1}{\mathrm{vol}_W(B_Y\cap W)^{1/k}}, \tag{4}$$

where G is a standard Gaussian random variable in $\mathbb{R}^m$. Banaszczyk's vector balancing theorem implies that $\alpha(\ell_2,Y) \lesssim \mathbb{E}\|G\|_Y$, and, together with the above estimate, this gives the tightness of (3) for Hilbert spaces.

The bounds (3) and (4) together also imply a partial converse to Banaszczyk's thereom: for any Banach space Y on $\mathbb{R}^m$,

$$\alpha(\ell_2^m,Y) \lesssim \mathbb{E}\|G\|_Y \lesssim (\log m)^2\alpha(\ell_2^m,Y).$$

Another approach to characterizing $\alpha(X,Y)$ geometrically was raised during working group discussions.

Conjecture 2. *For any Banach space X on $\mathbb{R}^m$, there exists a Hilbert space H, such that $\|v\|_H \le \|v\|_X$ for any $v \in \mathbb{R}^m$, and $\alpha(H,Y) \lesssim (\log m)^{O(1)}\alpha(X,Y)$.*

Note that the reverse inequality $\alpha(X,Y) \le \alpha(H,Y)$ is trivial. Then we can use the geometric characterization above for Hilbert space domains to characterize $\alpha(X,Y)$. Inequalities (1) and (2) can be seen to imply the existence of such a Hilbert space whenever $Y = \ell_p$, for $2 \le p \le \infty$. During a working group discussion, it was suggested that a possible counterexample to the existence of such a Hilbert space for every H is the case $X = Y = S_\infty^{n,n}$, where $S_\infty^{n,n}$ is the space of $n \times n$ real matrices, with the norm equal to the largest singular value (i.e. the Schatten ∞-norm, or the operator norm when we treat the matrices as operators from ℓ_2^n to ℓ_2^n). The conjecture was raised that the optimal choice of a Hilbert space in this setting is the one induced by the Hilbert-Schmidt norm, scaled down by a factor $\sqrt{n}$.

The Steinitz Constant Problem

A problem closely related to vector balancing is that of determining the Steinitz constant for a Banach space X on $\mathbb{R}^m$, defined as

$$\sigma_n(X) = \sup\left\{ \min_{\pi\in S^n} \max_{k=1}^{n} \Big\| \sum_{i=1}^{k} v_{\pi(i)} \Big\|_X : v_1,\dots,v_n \in \mathbb{R}^m, \max_{i=1}^{n} \|v_i\|_X \le 1, \sum_{i=1}^{n} v_i = 0 \right\};$$
$$\sigma(X) = \sup_{n\ge 1} \sigma_n(X)$$

where S_n is the set of permutations on $[n]$. Sevastyanov showed [Sevast-steinitz] (see also the survey [Barany-survey]) that $\sigma(X) \le m$, even when the norm X is not symmetric, i.e. we do not need to require $\|-v\|_X = \|v\|_X$ (but we do need homogeneity and the triangle inequality). This is tight up to constants for ℓ_1^m, and it's an open problem to show that $\sigma(\ell_2^m) = O(\sqrt{m})$ and $\sigma(\ell_\infty^m) = O(\sqrt{m})$. A non-constructive result of Banaszczyk [Bana13], building upon his vector balancing results, shows that $\sigma_n(\ell_2) = O(\sqrt{m} + \sqrt{\log n})$, which solves the problem for $n = 2^{O(m)}$. However, there is no efficient algorithm known that finds a permutation achieving this bound.

One of the working groups developed a polynomial time algorithm that, given vectors $v_1,\dots,v_n \in \mathbb{R}^m$ such that $\max_{i=1}^n \|v_i\|_2 \le 1$ and $\sum_{i=1}^n v_i = 0$, finds a permutation π for which $\max_{k=1}^n \Big\|\sum_{i=1}^k v_{\pi(i)}\Big\|_2 = O(\sqrt{m}\log n)$, as well as the analogous result for ℓ_∞^m. The algorithm uses a "splitting" technique based on combinatorial discrepancy, and discrepancy minimization algorithms.

A result of Chobanyan [Chobanyan94] shows that $\sigma_n(X)$ is bounded by the signed series constant

$$\rho_n(X) = \sup\left\{ \min_{\varepsilon_1,\dots,\varepsilon_n} \max_{k=1}^{n} \Big\| \sum_{i=1}^{k} v_i \Big\|_X : v_1,\dots,v_n \in \mathbb{R}^m, \max_{i=1}^{n} \|v_i\|_X \le 1 \right\}.$$

This motivates a computational question: given a sequence of vectors $v_1,\dots,v_n \in R^m$, approximate in polynomial time the optimal value of the minimization problem

$$\min_{\varepsilon_1,\dots,\varepsilon_n} \max_{k=1}^{n} \| \sum_{i=1}^{k} v_i \|_X.$$

In working group discussions, it was shown that the 1-dimensional case of this problem (in which all norms are equivalent) admits a polynomial-time approximation scheme via the classical rounding and dynamic programming approach. Little is known for the higher dimensional case of this problem, even for natural norms like $X \in \{\ell_2^m, \ell_\infty^m\}$. A lower bound on the minimum is given by $\max_{i=1}^n \|v_i\|_X$, and an upper bound which is bigger by a factor $\sqrt{m}(\log m)^{O(1)}$ follows from bounds on $\rho_n(X)$.

Compact List of Colorings for Hereditary Discrepancy

For a matrix A, the bound on hereditary discrepancy in (1) says that for any J, there is a coloring for A_J attaining the claimed bound. The coloring itself depends on J and it is natural to ask if the number or required colorings can be limited. Suppose that we define

$$\mathrm{herdisc}(A, K) := \min_{x^1,\dots,x^K \in \{-1,1\}^n} \max_{J \subseteq [n]} \min_{k \in [K]} |A_J x_J^k|_\infty$$

as the hereditary discrepancy of A with the restriction that only K many different colorings are allowed for coloring the induced set systems. Clearly, $\mathrm{herdisc}(A, 2^n - 1) = \mathrm{herdisc}(A)$. We asked

Question 1. *Is there a polynomial function $p(m, n)$ such that for every A, herdisc$(A, p(m, n)) \leq$ polylog$(m, n) \cdot \gamma_2(A)$.*

In working group discussions it was shown that the answer to Question 1 is solidly no, and that an exponential-sized list of colorings is needed. It was shown that there is a 0-1 matrix $A \in \mathbb{R}^{m\times n}$ (equivalently a set system) with $m = \sqrt{n}$ sets so that (a) $\gamma_2(A) = \mathrm{herdisc}(A) = 1$, and yet, (b) for $K = 2^{n^{\frac{1}{4}}}$, $\mathrm{herdisc}(A, K) \geq cn^{\frac{1}{4}}$.

The set system is rather simple: it is made up of $m = \sqrt{n}$ disjoint sets of size $\sqrt{n}$. To prove that a small list of colors is not sufficient, we argue that every fixed coloring incurs large discrepancy for all but a tiny fraction of random restrictions of A.

Discrepancy of Pseudo-lines

A famous conjecture, due to Beck and Fiala, states that the discrepancy of any set system of maximum degree at most t is bounded by $O(\sqrt{t})$. Here the degree of an element $i \in [n]$ equals the number of sets it belongs to. Spencer proved the conjecture in the special case when the set system is given by the set of lines in a finite projective plane. Finite projective planes have n sets and are t-regular and t-uniform, meaning that the degree of every element and the size of every set is equal to t. Moreover, every two lines intersect in exactly one point, and every two points determine a unique line. One of the working groups discussed a generalization of Spencer's proof to set systems of "pseudo-lines", defined as t-regular t-uniform set systems in which every two distinct sets have intersection of size one. A discrepancy upper bound of $O(\sqrt{t})$ was proved when $t = n^\delta$ for a constant $\delta > 0$, but proving this result when $t = n^{o(1)}$ remains open. Note that in finite projective planes $t = \Theta(\sqrt{n})$.

Constructive Proof of Banaszczyk's Theorem

An open computational problem in combinatorial discrepancy theory is to give a polynomial time algorithm that, given vectors $v_1, \dots, v_n \in \mathbb{R}^m$, all of ℓ_2 norm at most 1, and a polytope K of Gaussian measure at least $1/2$, finds a sequence of signs $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$ such that $\sum_{i=1}^n \varepsilon_i v_i \in O(1) \cdot K$. Such an algorithm would give an algorithmic version of Banaszczyk's vector balancing theorem, and would also give a constructive version of (1), i.e. would give an efficient algorithm that constructs colorings achieving the bound in (1) for any submatrix of A. This problem was solved when K is a scaled m-dimensional cube in a recent paper by Bansal, Dadush, and Garg. This is an interesting special case, as it corresponds to the best bound known for the Komlos conjecture, which states that $\alpha(\ell_2^m, \ell_\infty^m) = O(1)$. Some discussions that led to these results happened in working groups during the workshop.

Conclusion

The workshop was successful in several of its goals. It brought together researchers from theoretical computer science, functional analysis, and geometric discrepancy theory. It identified a list of open problems in combinatorial discrepancy theory and related fields. It also led to progress on several of these problems, in terms of new approaches to their solution and in terms of partial results. The collaborative nature of an AIM workshop was instrumental in these successes.