B. News Scan 1999–2009

Geometric discrepancy is a lively field and many things have happened since the first appearance of this book ten years ago. In the present revised printing, scheduled to appear in 2009 or 2010, I decided to add this appendix mentioning some of the new results, rather than trying to insert dispersed remarks into the old text.

I should perhaps begin with a disclaimer. I have been following the development in discrepancy theory only cursorily, devoting most of my time to other subjects. The following remarks should not be regarded as a serious survey. Among the results I happened to learn about, I've selected according to strictly objective scientific criteria: the results I liked best, those I considered interesting, unexpected, or particularly difficult, those easy to write about, those proved by me or my friends, and so on.

Boxes in dimensions 3 and more. The closest to the heart of a classical discrepancy theorist are probably two recent papers improving lower bounds on $D(n, \mathcal{R}_d)$, the Lebesgue-measure discrepancy for axis-parallel boxes.

We recall that Roth's lower bound for the L_2 average discrepancy gives $D(n, \mathcal{R}_d) = \Omega((\log n)^{(d-1)/2})$ for every fixed $d \geq 2$. A common belief, supported by a proof only for $d = 2$, is that the order of $D(n, \mathcal{R}_d)$ is at least by the factor of $\sqrt{\log n}$ larger. For many years the only step in this direction for $d \geq 3$ had been Beck's [Bec89c] magnificent proof improving Roth's bound in dimension 3 by the factor of roughly $(\log \log n)^{1/8}$.

In 2006 Bilyk and Lacey [BL08] simplified and greatly developed Beck's approach, improving the 3-dimensional lower bound to $\Omega((\log n)^{1+\eta})$ for a small constant $\eta > 0$ (which they didn't compute explicitly). Similar to Beck's proof, the the core of their method is a so-called *small ball inequality*, an inequality for multidimensional Haar functions i.e., higher-dimensional analogs of the functions f_j from Halász's proof (see Section 6.2 and its Exercise 2).

To state the inequality, let $\mathbf{r} = (r_1, \ldots, r_d)$ be a d-dimensional vector of nonnegative integers, let us write $|\mathbf{r}| = r_1 + \cdots + r_d$, and let $R_{\mathbf{r}}$ be the appropriate Rademacher function, given by $R_{\mathbf{r}}(x) = \prod_{i=1}^{d} (-1)^{\lfloor 2^{r_i+1} x_i \rfloor}$. A *weighted* $\mathbf{r}$*-function* is a function $f\colon [0,1]^d \to \mathbf{R}$ such that on every binary canonical box B of size $2^{-r_1} \times 2^{-r_2} \times \cdots \times 2^{-r_d}$, the function f coincides with $\alpha_B R_{\mathbf{r}}$ for some real α_B (depending on the box B). (The $\mathbf{r}$-function defined in Exercise 6.1.1 is a special case with $\alpha_B \in \{-1, 0, +1\}$ for all B.) In the small ball inequality we seek, for given natural numbers d and k, the smallest $C = C_{d,k}$ such that for every choice of weighted $\mathbf{r}$-functions $f_{\mathbf{r}}$, for all $\mathbf{r}$ with $|\mathbf{r}| \geq k$, we have

$$\sum_{\mathbf{r}:|\mathbf{r}|=k} \|f_{\mathbf{r}}\|_1 \leq C \bigg\| \sum_{\mathbf{r}:|\mathbf{r}|\geq k} f_{\mathbf{r}} \bigg\|_\infty .$$

A Roth-like L_2 averaging argument shows that $C = O(k^{(d-1)/2})$ for every fixed d, the *small ball conjecture* asserts $C = O(k^{(d-2)/2})$ for all $d \geq 2$ (which is known only for $d = 2$), and [BL08] proved $C_{3,k} = O(k^{1-\eta})$. The small ball inequality is of fundamental nature and it has applications in other fields (probability theory, approximation theory) as well. In particular, the name "small ball" comes from a probabilistic setting, concerning the behavior of the d-dimensional Brownian random walk.

The paper [BL08], available on ArXiv, uses lots of beautiful mathematics, mostly harmonic analysis (e.g., the Littlewood–Paley theory), and it is written in a way that looks quite accessible even to us non-experts in this field. Later Bilyk, Lacey, and Vagharshakyan [BLV08] extended the method to higher dimensions, obtaining $D(n, \mathcal{R}_d) = \Omega((\log n)^{(d-1)/2+\eta})$ for every fixed d and some positive $\eta = \eta(d)$, again through the corresponding small ball inequality.

The discrepancy function for corners in the plane. Bilyk, Lacey, Parissis, and Vagharshakyan [BLPV08] improved our understanding of the discrepancy function for two-dimensional corners. We recall that $D(n, \mathcal{C}_2)$, the worst-case, or L_∞, discrepancy for corners, is of order $\log n$, while the L_p average discrepancy $D_p(n, \mathcal{C}_2)$ is of order $\sqrt{\log n}$ for every fixed $p \in [1, \infty)$.

Bilyk et al. proved bounds that, in a sense, smoothly interpolate between these two results: they obtained a tight bound, of order $(\log n)^{1-1/\alpha}$, for the Orlicz norms $\|.\|_{\exp(L^\alpha)}$ of the discrepancy function, for every fixed $\alpha \in [2, \infty)$. We recall that the Orlicz norm is a generalization of the L_p norm where the numeric parameter p is replaced with a (convex) real *function* ψ. The Orlicz norm of a function f (defined on a space X with measure μ) equals $\inf\{t > 0\colon \int_X \psi(|f(x)|/t)\,\mathrm{d}\mu(x) \leq 1\}$; the L_p norm is recovered for $\psi(x) = |x|^p$. In the result cited above we have $\psi(x) = e^{|x|^\alpha}$, which means that the norm is even much more influenced by large fluctuations that the L_p norms and thus it is a "closer approximation" of the L_∞ norm.

Explicit constructions for L_p discrepancy. Chen and Skriganov [CS02] obtained an explicit construction of a set meeting Roth's lower bound for the L_2 discrepancy for corners, in every fixed dimension (while all of the several constructions known before had some probabilistic component); also see [CS08] for a substantial simplification of the proof. We won't describe the construction here; we just mention that it has some features in common with the construction of b-ary nets in Section 2.3, dealing with a suitable vector subspace of $GF(b)^{md}$ (for a prime b) and then mapping it to a point set in $[0, 1]^d$ in the usual way, by reading the components as digits in base b. Skriganov [Skr06] constructed explicit sets in the unit cube with asymptotically optimal L_p discrepancy for every fixed $p \in (1, \infty)$ and every fixed dimension d.

Extra-large discrepancy for hyperbolic needles. Beck [Beca], [Becb] investigated, in our language, the discrepancy for translated and rotated copies of the *hyperbolic needle* $H_\gamma(n) = \{(x, y) \in \mathbf{R}^2\colon x \in [1, n], |xy| \le \gamma\}$. We note that the area $\mathrm{vol}(H_\gamma(n)) = 2\gamma \ln n$. The number of *integer points* in such rotated and translated hyperbolic needle (essentially) corresponds to the number of integer solutions (x, y) with $x \ge 1$, $1 \le y \le n$ of the *inhomogeneous Pell equation* $|(x+\beta)^2 - \alpha y^2| \le \gamma$, which is a quantity of considerable interest in number theory.

Beck established an "extra-large discrepancy" phenomenon. If P is the integer lattice $\mathbf{Z}^2$ or, more generally, a set in $\mathbf{R}^2$ of density 1 in which every two points have distance at least σ (a positive constant), then for 99 percent of rotational angles θ, there is a translated copy H of $H_\gamma(n)$ rotated by θ such that $|P \cap H|$ differs from $\mathrm{vol}(H)$ by $\Omega(\log n)$, i.e., by a *fixed fraction* of the area, the constant depending on γ and σ. (We gloss over some subtleties of Beck's result; see his Theorem 4 for a stronger formulation.)

Now let $\gamma > 0$ be fixed and, for $\beta \in [0, 1]$, let $\tilde{H}^\beta_\gamma(n)$ be $H_\gamma(n)$ rotated by 45 degrees and translated by β in the positive x-direction. We set $F_n(\beta) := |\mathbf{Z}^2 \cap \tilde{H}^\beta_\gamma(n)|$. Beck [Becb] discovered that, for $\beta \in [0, 1]$ chosen uniformly at random, the distribution of $F_n(\beta)$ suitably normalized tends to the standard normal distribution (and in particular, the "typical" discrepancy of $\tilde{H}^\beta_\gamma(n)$ is of order $\sqrt{\log n}$). Moreover, $F_n(\beta)$ also satisfies a law of the iterated logarithm.

L_1 discrepancy for halfspaces and lattice points in polyhedra. Chen and Travaglini [CT09b] extended Proposition 3.4 to an arbitrary dimension, showing that the L_1 discrepancy for halfspaces in $\mathbf{R}^d$ is at most $O(\log^d n)$, attained for appropriately re-scaled $\mathbf{Z}^d$. The proof is based on results of Brandolini, Colzani, and Travaglini [BCT97] (plus some "boundary effects" have to be dealt with). In the latter paper it was proved, among others, that if C is a fixed polyhedron in $\mathbf{R}^d$ (not necessarily convex), then the expected discrepancy of a randomly rotated and translated copy of C w.r.t. the lattice $\frac{1}{m}\mathbf{Z}^d$ is bounded by $O(\log^d m)$.

The main theme of [BCT97] is the "average decay" of a Fourier transform, a more or less classical topic. Letting C be a compact set in $\mathbf{R}^d$, one studies the behavior of $\widehat{\chi_C}$, the Fourier transform of the characteristic function of C. In particular, in the setting of [BCT97], one takes some L_p average of $\widehat{\chi_C}$ over the sphere of radius R and investigates how fast it tends to 0 as $R \to \infty$. This is highly relevant for discrepancy lower bounds in the style of Chapter 7, as well as for questions about lattice point distributions in copies of C; see Travaglini [Tra04] for a nice survey.

More on lattice points. The last few results mentioned above are relevant for geometric discrepancy, but they really belong to the geometry of numbers or, more precisely, theory of irregularities of distribution for the integer lattice $\mathbf{Z}^d$. This is an extensive area on its own, of much more number-theoretic nature than discrepancy theory in general, and with deep connections to harmonic analysis and other fields. Here we mention two interesting discrepancy-related topics.

Let $p = p(x_1, \ldots, p_d)$ be a d-variate polynomial with integer coefficients. A fundamental problem in number theory is to find integer solutions of $p(x) = \lambda$, where $\lambda \in \mathbf{Z}$. Geometrically, one looks for integer points on the level surface $\{x \in \mathbf{R}^d\colon p(x) = \lambda\}$. Magyar [Mag07] studied the equidistribution of these point sets for the case of p positive and homogeneous, and in particular, their discrepancy for caps (i.e., intersections of the level surface with halfspaces). Among other amazing results he proved, that for $p(x) = x_1^2 + \cdots + x_d^2$, where the level surface is a sphere, these sets have an almost optimal discrepancy, up to an n^ε factor (among all possible sets of the same size in the sphere), for almost all caps. Roughly speaking, the exceptional caps not covered by this bound have normal directions that are "too well approximable" by rational directions.

The next topic concerns the L_2 discrepancy for balls. For definiteness, let us consider the toroidal discrepancy; see the notes to Section 7.1. Let P be a fixed n-point set in the unit torus $T^d = \mathbf{R}^d/\mathbf{Z}^d$, let $r \in (0, \frac{1}{2})$ be a given radius, and let $D_2(r)$ denote the L_2 average of the discrepancy of a ball of radius r centered at x, averaged over x uniformly distributed in T^d. Results of Beck and of Montgomery (see [BC87], [Mon94]) show that the average of $D_2(r)$ over $r \in (0, \frac{1}{2})$ is at least of order $n^{1/2-1/2d}$.

Now let the set P be the scaled grid $\frac{1}{m}\mathbf{Z}^d$, with an integer m; this is an n-point set in T^d, $n = m^d$. It is known that this P matches, up to a constant factor, the just mentioned lower bound (for the average over r). However, a surprising phenomenon, discovered by Parnovski and Sobolev [PS01] (Section 3), appears when one considers $D_2(r)$ for $r \in (0, \frac{1}{2})$ *fixed*. The behavior depends on the remainder of the dimension d modulo 4: for $d \not\equiv 1 \pmod 4$, $D_2(r)$ behaves "regularly", being always of order $n^{1/2-1/2d}$, but for $d \equiv 1 \pmod 4$ there are infinitely values of m for which $D_2(r)$ is asymptotically *smaller*, namely, of order at most $n^{1/2-1/2d}(\log n)^{-c_d}$ (with an explicit constant $c_d > 0$). From below Parnovski and Sobolev proved $D_2(r) = \Omega(n^{1/2-1/2d-\delta})$ for every fixed $\delta > 0$; Konyagin, Skriganov, and Sobolev [KSS03] improved this, replacing $n^{-\delta}$ by $e^{-O((\log\log n)^4))}$.

This phenomenon plays a significant role in Chen and Travaglini [CT09a], who also considered the L_2 toroidal discrepancy for balls and whose goal was comparing a deterministic construction, namely, the scaled grid as above, with a randomized construction in the spirit of "jittered sampling", where one starts with the grid points and randomly perturbs each of them independently of the others. They found that the grid is better in small dimensions, while the randomized construction wins in large dimensions, *except* for dimensions $d \equiv 1 \pmod 4$, where the grid is better for infinitely many values of m due to the Parnovski–Sobolev result. Similar investigations in a more general setting were undertaken by Brandolini et al. [BCGT09].

Discrepancy for high-dimensional corners. An interesting question is, how $D(n, \mathcal{C}_d)$, the (worst-case) discrepancy for corners, behaves for d large, say comparable to n? In particular, Heinrich et al. [HNWW01] investigate the quantity $n_\infty(d, \varepsilon) = \min\{n\colon D(n, \mathcal{C}_d)/n \le \varepsilon\}$; that is, the smallest number of points in $[0,1]^d$ that can approximate the measure of all corners with *relative* accuracy ε. Perhaps surprisingly, $n_\infty(d, \varepsilon)$ is *polynomially bounded* in d and $\frac{1}{\varepsilon}$. (This should be contrasted with the fact that for $d = \log_2 n$, say, we have $D(n, \mathcal{C}_d) = 2^{\Omega(d)}$, as can be calculated from Roth's lower bound—see, e.g., [Mat98b] for the appropriate formulas.) Indeed, a straightforward VC-dimension argument yields $n_\infty(d, \varepsilon) \le Cd\varepsilon^{-2} \log \frac{d}{\varepsilon}$, with an explicit constant C (independent of d, of course!), and using a deep result of Talagrand, this can be improved to $Cd\varepsilon^{-2}$—see [HNWW01].

The best known lower bound is due to Hinrichs [Hin04]: $n_\infty(d, \varepsilon) \ge cd/\varepsilon$, for some constant $c > 0$, all $\varepsilon > 0$ smaller than a suitable constant, and all d. The idea of this lower bound is simple. One constructs a large set $\mathcal{N}_\varepsilon \subset \mathcal{C}_d$ of corners such that the symmetric difference of every two has volume exceeding ε. If P is an n-point set with discrepancy at most εn, then $P \cap C \ne P \cap C'$ for every two corners $C \ne C'$ in $\mathcal{N}_\varepsilon$. Finally, the number of different intersections of P with corners is estimated using a VC-dimension argument.

The cited polynomial upper bounds are probabilistic—they hold for a typical random n-point set. An interesting open problem is obtaining an *explicit* construction of polynomial size. What is meant by "explicit"? This word is often used in an informal sense, but theoretical computer science offers a formal definition: explicit means computable by a deterministic polynomial-time algorithm, in our case in time polynomial in d and $\frac{1}{\varepsilon}$. Methods of theoretical computer science, developed mainly for the purpose of derandomizing probabilistic algorithms, have also led to the strongest results so far. Namely, the work of Even et al. [EGL+92] provides explicit sets witnessing $n_\infty(d, \varepsilon) \le (d/\varepsilon)^{O(\log d)}$, and also $n_\infty(d, \varepsilon) \le (d/\varepsilon)^{O(\log(1/\varepsilon))}$ (which is polynomial in d for ε fixed).[1] The second bound has later been improved; to my knowledge, the best result is $n_\infty(d, \varepsilon) \le d^{O(1)} \varepsilon^{-O(\sqrt{\log(1/\varepsilon)})}$ following from Lu [Lu02]. All of these constructions are actually formulated for the discrete grid; that is, instead of the Lebesgue measure on $[0,1]^d$ one approximates the counting measure on the grid $\{1, 2, \ldots, q\}^d$ (for converting this to the Lebesgue-measure case, one needs to set $q = Cd/\varepsilon$). The constructions work not only for corners, but also for *combinatorial rectangles*; see the notes on page 34.

[1] In contrast, the bounds known for the usual constructions for fixed d, such as the Halton–Hammersley sets, have at least exponential dependence on d.

There are also nontrivial results concerning deterministic computation of sets witnessing $n_\infty(d,\varepsilon) = O(d\varepsilon^{-2}\log\frac{1}{\varepsilon})$, almost matching the best known probabilistic bound, but the running time of these algorithms are exponential in d; see, e.g., Doerr and Gnewuch [DG08].

The trace bound. An interesting lower bound technique for combinatorial discrepancy, the so-called *trace bound*, was developed by Chazelle and Lvov [CL01], which, for example, yields direct proofs for some results where previously one had to go via the Lebesgue-measure discrepancy. It asserts that, for a set system $\mathcal{S}$ on n points, with *at most* n sets, and with incidence matrix A, we have

$$\operatorname{disc}(\mathcal{S}) \geq \frac{1}{4} \cdot 324^{-n\cdot \operatorname{tr}((A^TA)^2)/t^2} \sqrt{t/n},$$

where $t = \operatorname{tr}(A^TA)$ and $\operatorname{tr}(M)$ denotes the trace (sum of diagonal elements) of a matrix M.

Adding a single set. A tantalizing open question in combinatorial discrepancy is, by how much can the hereditary discrepancy of a set system on n points increase by adding a single set? The truth could perhaps be an additive constant, but the current best result of Kim, Matoušek, and Vu [KMV05] gives only a *multiplicative factor* of $O(\log n)$, with a half-page proof.

Linear discrepancy versus hereditary discrepancy. We have seen that the linear discrepancy of any set system, or more generally, of any matrix, is no more than twice the hereditary discrepancy. Spencer conjectured that the factor 2 can be improved to $2(1-\frac{1}{n+1})$ for all matrices with n columns (which, if true, is tight). Doerr [Doe04a] and, later but independently, Bohman and Holzman [BH05] proved the special case of this conjecture with A totally unimodular. Both proofs are nice and the second one is also quite short.

Multicolor discrepancy. The notion of combinatorial discrepancy has been generalized from two colors to k colors. That is, we want to color the ground set with k colors so that each set has roughly $\frac{1}{k}$ fraction of each color; see Doerr and Srivastav [DS03] for a survey. While many of the results are direct generalizations from the 2-color case, some interesting phenomena have been found. In particular, Doerr [Doe04b] showed, with a neat proof employing the k-color *linear* discrepancy, that the hereditary discrepancy of a set system $\mathcal{S}$ is nearly independent of the number of colors; that is, for every $k, \ell \geq 2$ there is a constant $C = C(k,\ell)$ such that the ℓ-color hereditary discrepancy of $\mathcal{S}$ is at most C-times the k-color hereditary discrepancy. On the practical side, multicolor discrepancy turned out to be important in a problem of storing data on parallel disks, as was observed independently by Chen and Cheng [CC04] and by Doerr, Hebbinghaus, and Werth [DHW06].