10 THE DISCREPANCY METHOD IN COMPUTATIONAL GEOMETRY

B. Chazelle

INTRODUCTION

Discrepancy theory investigates how uniform nonrandom structures can be. For example, given n points in the plane, how should we color them red and blue so as to minimize the difference between the number of red points and the number of blue ones within any disk? Or, how should we place n points in the unit square so that the number of points that lie within any given triangle in the square is as close as possible to n times the area of the triangle? Questions of this nature have direct relevance to computational geometry for two reasons. One of them is their close association with the problem of derandomizing probabilistic algorithms. Such algorithms are often based on random sampling and discrepancy theory provides tools for carrying out the sampling deterministically. This has led to the intriguing fact that virtually all of the important problems in low-dimensional computational geometry can be solved as efficiently deterministically as probabilistically. The second application of discrepancy theory to computational geometry is in the area of lower bounds for multidimensional searching. The complexity of these problems is often tied to spectral properties of geometric set systems, which themselves lie at the heart of geometric discrepancy theory.

10.1 VC-DIMENSION THEORY

GLOSSARY

Set system: A pair $\Sigma = (X, \mathcal{R})$, where X is a set and $\mathcal{R}$ is a collection of subsets of X, is called a set system. The terminology *geometric set system* refers to the case where $X \subset \mathbf{R}^d$ and each $R \in \mathcal{R}$ is of the form $X \cap f(C)$, where C is a fixed region of $\mathbf{R}^d$ (eg, a simplex) and f is any member of a fixed group F of transformations (eg, a rotation).

VC-dimension: Given $Y \subseteq X$, the set system *induced* by Y is of the form $(Y, \mathcal{R}|_Y)$, where $\mathcal{R}|_Y = \{\, Y \cap R \mid R \in \mathcal{R} \,\}$. The *VC-dimension* of Σ is the maximum size of any Y such that $\mathcal{S}|_Y = 2^Y$. For example, the VC-dimension of the infinite set system formed by points in $\mathbf{R}^2$ and halfplanes is 3. The *shatter function* $\pi_{\mathcal{R}}(m)$ of a (usually infinite) set system $\Sigma = (X, \mathcal{R})$ is the maximum number of subsets in the set system $(Y, \mathcal{R}|_Y)$ induced by any $Y \subseteq X$ of size m. If $\pi_{\mathcal{R}}(m)$ is bounded by cm^d, for some constants $c, d > 0$, then the set system is said to have a *shatter function exponent* of at most d.

0-8493-8524-5/97/$0.00+$.50

Dual set system: The set system $\Sigma^* = (X^*, \mathcal{R}^*)$, where $X^* = \mathcal{R}$, $\mathcal{R}^* = \{ R_x \mid x \in X \}$, and $R_x = \{ R \in \mathcal{R} \mid x \in R \}$, is called the *dual set system* of Σ. The shatter function of Σ^* is called the *dual shatter function* of Σ.

Discrepancy: The *incidence matrix* A of a finite set system $\Sigma = (X, \mathcal{R})$ is the matrix whose $|X|$ columns (resp. $|\mathcal{R}|$ rows) are indexed by the elements of X (resp. $\mathcal{R}$): A_{ij} is 1 if the i-th set of $\mathcal{R}$ contains the j-th element of X, and 0 otherwise. The (red-blue) *discrepancy* of Σ is $\min_{x \in \{-1,1\}^{|X|}} \|Ax\|_\infty$.

The concept of VC-dimension was introduced by Vapnik and Chervonenkis [VC71]. The relation between VC-dimension and shatter function is a key component of the theory.

LEMMA 10.1.1 [VC71, Sau72, She72]

If the shatter function exponent is $O(1)$, then so is the VC-dimension. Conversely, if the VC-dimension is $d \geq 1$ then, for any $m \geq d$, $\pi_{\mathcal{R}(m)} < (em/d)^d$.

LEMMA 10.1.2 [Ass83]

If a set system has VC-dimension d, then its dual has VC-dimension less than 2^{d+1}.

Any set system of n elements and n sets has discrepancy $O(\sqrt{n})$, and this bound is sometimes tight. If the VC-dimension is bounded, the discrepancy falls below the $\sqrt{n}$ barrier. The bounds below are stated in terms of the shatter function exponent. In view of Lemma 10.1.1, we can replace the exponent by the VC-dimension if we wish. Matoušek, Welzl, and Wernisch [MWW93] established a bound of $O(n^{1/2-1/2d}(\log n)^{1+1/2d})$ on the discrepancy of set systems with shatter function exponent d. This was improved to $O(n^{1/2-1/2d})$ by Matoušek:

THEOREM 10.1.3 [Mat95b]

The discrepancy of a set system of n elements with shatter function exponent $d > 1$ is $O(n^{1/2-1/2d})$, which is optimal for $d \geq 2$.

Similar bounds can be obtained in terms of the dual shatter function. Matoušek, Welzl, and Wernisch [MWW93] proved a bound of $O(n^{1/2-1/2d}\sqrt{\log n})$ on the discrepancy of set systems with dual shatter function exponent d. It is surprising that an extra $\sqrt{\log n}$ should be needed. Optimality was shown by Matoušek [Mat97] for the cases $d = 2, 3$, and by Alon, Rónyai, and Szabó [ARS99] for $d > 3$.

THEOREM 10.1.4 [MWW93]

The discrepancy of a set system of n elements with dual shatter function exponent $d > 1$ is $O(n^{1/2-1/2d}\sqrt{\log n})$, which is optimal for $d \geq 2$.

10.2 SAMPLING IN BOUNDED VC-DIMENSION

GLOSSARY

ε-Net: Given a finite set system $(X, \mathcal{R})$ and any $0 < \varepsilon < 1$, a set $N \subseteq X$ is called an *ε-net for* $(X, \mathcal{R})$ if $N \cap R \neq \emptyset$, for any $R \in \mathcal{R}$ with $|R|/|X| > \varepsilon$.

ε-Approximation: Given a finite set system $(X, \mathcal{R})$ and any $0 < \varepsilon < 1$, a set $A \subseteq X$ is called an *ε-approximation for* $(X, \mathcal{R})$ if, for any $R \in \mathcal{R}$,

$$\left| \frac{|R|}{|X|} - \frac{|A \cap R|}{|A|} \right| \leq \varepsilon.$$

Product set system: Given two finite set systems $\Sigma_1 = (X_1, \mathcal{R}_1)$ and $\Sigma_2 = (X_2, \mathcal{R}_2)$, the *product set system* $\Sigma_1 \otimes \Sigma_2$ is defined as $(X_1 \times X_2, \mathcal{T})$, where $\mathcal{T}$ consists of all subsets $T \subseteq X_1 \times X_2$ such that each set of the form $T^1_{x_2} = \{\, x \in X_1 \,|\, (x, x_2) \in T \,\}$ belongs to $\mathcal{R}_1$ and, similarly, $T^2_{x_1} = \{\, x \in X_2 \,|\, (x_1, x) \in T \,\}$ belongs to $\mathcal{R}_2$.

To sample a set system Σ is to extract a (small) subset of the elements whose intersection with any set R of Σ is a good predictor of the size of R. This is the idea behing an ε-approximation. A weaker version of sampling, the ε-net, requires only that large enough sets R be intersected by the sample. The key result in VC-dimension theory is that if Σ has bounded VC-dimension, then for any given level of accuracy, the sample size need not depend on the size of the set system. This is rather counterintuitive. It says, for example, that if we want to estimate how many people live within 1 mile of a post office and, to go about it, we opt to pick a sample of the population, solve the problem for the sample, and then scale up the answer appropriately, then the same sample size works just as well whether the country is France or India!

LEMMA 10.2.1

Let X_1, X_2 be disjoint subsets of X of the same size, and let A_i be an ε-approximation for the subsystem induced by X_i. If $|A_1| = |A_2|$, then $A_1 \cup A_2$ is an ε-approximation for the subsystem induced by $X_1 \cup X_2$.

LEMMA 10.2.2

If A is an ε-approximation for $(X, \mathcal{R})$, then any ε'-approximation (resp. -net) for $(A, \mathcal{R}|_A)$ is also an $(\varepsilon + \varepsilon')$-approximation (resp. -net) for $(X, \mathcal{R})$.

A greedy approach to sampling yields an effective algorithm for arbitrary set systems. Writing $\varepsilon = 1/r$, choose some $1 \leq r \leq n$. First, remove all sets $R \in \mathcal{R}$ of size at most n/r. Second, initialize the set N to $\emptyset$. Next, find the element $x \in X$

that belongs to the most sets of $\mathcal{R}$ (in case of a tie, any one will do) and add it to N. Remove from $\mathcal{R}$ every set that contains x, discard x, and iterate in this fashion until $\mathcal{R}$ is empty. An elementary analysis shows that this produces a $(1/r)$-net for $(X, \mathcal{R})$ of size $O(r \log |\mathcal{R}|)$. This was proven independently by Johnson [Joh74] and Lovász [Lov75]. A slightly more complicated, "weigted" version of the greedy algorithm, due to Chazelle [Cha00], gives an analog result for $(1/r)$-approximations.

THEOREM 10.2.3

Given a set system $(X, \mathcal{R})$, where $|X| = n$ and $|\mathcal{R}| = m$, for any $1 \le r \le n$, it is possible to find, in time $O(nm)$, a $(1/r)$-net for $(X, \mathcal{R})$ of size $O(r \log m)$ and a $(1/r)$-approximation for $(X, \mathcal{R})$ of size $O(r^2 \log m)$.

The size of the sample depends (albeit weakly) on the size of the set system. In the presence of bounded VC-dimension, however, this dependency magically disappears. Again, we will base our results not on the VC-dimension but on the shatter function exponent d (but the same results hold if d denotes the VC-dimension). Geometric set systems often are defined implicitly and are accessible via an *oracle* function that takes any $Y \subseteq X$ as input and returns the list of sets in $\mathcal{R}|_Y$ (each set represented explicitly). We assume that the time to complete this task is $O(|Y|^{d+1})$, which is linear in the maximum possible size of the oracle's output. The existence of such an oracle is quite realistic: For example, in the case of points and disks in the plane, we have $d = 3$, and so this assumes that, given n points, we can enumerate all subsets enclosed by a disk in time $O(n^4)$. To do this, enumerate all k-tuples of points ($k \le 3$) and, for each tuple, find which points lie inside the smallest disk enclosing the k points.

THEOREM 10.2.4

Given a set system $(X, \mathcal{R})$ of shatter function exponent d, for any any $r \ge 2$, a $(1/r)$-approximation for $(X, \mathcal{R})$ of size $O(dr^2 \log dr)$ and a $(1/r)$-net for $(X, \mathcal{R})$ of size $O(dr \log dr)$ can be computed in time $O(d)^{3d}(r^2 \log dr)^d |X|$.

A randomized construction of ε-approximations in bounded VC-dimension was given by Vapnik and Chervonenkis [VC71]. The deterministic construction cited above is due to Chazelle and Matoušek [CM96]. Earlier influential work can be found in [CF90, Mat90, Mat91, Mat95a]. The bound on the size of ε-nets was established by Haussler and Welzl [HW87]. The running time for computing a $(1/r)$-net was improved to $O(d)^{3d}(r \log dr)^d |X|$ by Brönnimann, Chazelle, and Matoušek [BCM99], using the concept of a *sensitive ε-approximation*.

For fixed d, Komlós, Pach, and Woeginger [KPW92] showed that the bound of $O(r \log r)$ for $(1/r)$-nets cannot be improved in general (see a nice discussion in [PA95]). The situation is different with ε-approximations, however, for which Theorems 10.1.3 and 10.1.4 can be put to use. Matoušek, Welzl, and Wernisch proved the following:

THEOREM 10.2.5 [MWW93]

Let $(X, \mathcal{R})$ be a set system of VC-dimension $d > 1$. There exists a $(1/r)$-approximation for $(X, \mathcal{R})$ of size $O(r^{2-2/(d+1)}(\log r)^{2-1/(d+1)})$, for any $r \ge 2$.

The log factor can be removed by appealing to Theorem 10.1.3.

THEOREM 10.2.6 [MWW93]

Let $(X, \mathcal{R})$ be a set system with dual shatter function exponent $d > 1$. There exists a $(1/r)$-approximation for $(X, \mathcal{R})$ of size $O(r^{2-2/(d+1)}(\log r)^{1-1/(d+1)})$, for any $r \geq 2$.

Given n lines in the plane, we can use an ε-approximation to estimate how many lines cut through an arbitrary line segment. Suppose that, instead, we wish to estimate the number of vertices in the induced arrangement that fall within an arbitrary triangle. Product set systems allow us to do that. Let Σ_1 be the set system induced by n blue lines in the plane and the set of all line segments: a set of the system is the subset of blue lines intersected by a given segment. We define Σ_2 similarly with n red lines. The product $\Sigma_1 \otimes \Sigma_2$ is a set system $(Z, \mathcal{T})$, where Z is the set of red-blue vertices of the induced arrangement (assuming general position). A set of $\Sigma_1 \otimes \Sigma_2$ is any subset T of Z such that, along any (blue or red) line ℓ, the vertices of T incident to ℓ (if any) appear consecutively among the red-blue vertices of ℓ. This suggests we can use ε-approximations to, say, estimate how many red-blue vertices fall in an arbitrary triangle, or even an arbitrary convex region. One must be careful, however. The product of two set systems with bounded VC-dimension might not itself have bounded VC-dimension. Indeed, any bichromatic matching of the lines gives a collection of n vertices, any of whose 2^n subsets is a valid set of $\mathcal{T}$. Although the product does not have bounded VC-dimension, it so happens that sampling in it is still possible: that is the beauty of product set systems.

LEMMA 10.2.7

Given any $0 \leq \varepsilon_i \leq 1$, let A_i be an ε_i-approximation for a set system Σ_i, for $i = 1, 2$. Then, the Cartesian product $A_1 \times A_2$ is an $(\varepsilon_1 + \varepsilon_2)$-approximation for $\Sigma_1 \otimes \Sigma_2$.

The product operation is associative, and the theorem can be extended to multiple products of set systems. The notion of product set system was introduced by Brönnimann, Chazelle, and Matoušek, who also proved:

LEMMA 10.2.8 [BCM99]

Given an ε-approximation A for a set system Σ, the d-fold Cartesian product $A \times \cdots \times A$ is a $(d\varepsilon)$-approximation for the d-fold product $\Sigma \otimes \cdots \otimes \Sigma$.

One of the most important applications of the theorem above is to counting vertices in an arrangement of hyperplanes in $\mathbf{R}^d$. We consider the set system $\Sigma = (H, \mathcal{R})$ formed by a set H of hyperplanes in $\mathbf{R}^d$, where each $R \in \mathcal{R}$ is the subset of H intersected by an arbitrary line segment. Given a convex body σ (not necessarily full-dimensional), consider the arrangement formed by H within the affine span of σ, ie, the lowest-dimensional flat that contains σ, and let $V(H, \sigma)$ be the set of vertices of this arrangement that lie inside σ.

THEOREM 10.2.9 [Cha93, BCM99]

Given a set H of hyperplanes in $\mathbf{R}^d$ in general position, along with an ε-approximation A for $\Sigma = (H, \mathcal{R})$, for any convex body σ of dimension $k \leq d$,

$$\left| \frac{|V(H,\sigma)|}{|H|^k} - \frac{|V(A,\sigma)|}{|A|^k} \right| \leq \varepsilon.$$

10.3 GEOMETRIC ALGORITHMS

GLOSSARY

ε-Cutting: Given a set H of n hyperplanes in $\mathbf{R}^d$ and $\varepsilon > 0$, a collection $\mathcal{C}$ of closed full-dimensional simplices (some of them unbounded) is called an ε-cutting if: (i) their interiors are pairwise disjoint, and together they cover $\mathbf{R}^d$; (ii) the interior of any simplex of $\mathcal{C}$ is intersected by at most εn hyperplanes of H.

Simplicial partition: Given a finite set $P \subset \mathbf{R}^d$, a collection $\{(P_i, R_i)\}$ is a *simplicial partition*, if (i) the P_i's partition P and (ii) each R_i is a relatively open simplex enclosing P_i. The R_i's can be of any dimension and need not be disjoint, and P_i need not be equal to $P \cap R_i$. We say that a hyperplane *cuts* R_i if it intersects, but does not contain, R_i. The maximum number of R_i's that a single hyperplane can cut is the *cutting number* of the simplicial partition.

Partition tree: Given a finite set $P \subset \mathbf{R}^d$, a *partition tree* for P is a rooted tree $\mathcal{T}$ whose root is associated with the point set P. The set P is partitioned into subsets $P_1, \ldots, P_m$, and each P_i is associated with a distinct child v_i of the root. There is a convex open set R_i, called the *region* of v_i, that contains P_i. The regions R_i are not necessarily disjoint. If $|P_i| > 1$, the subtree rooted at v_i is defined recursively with respect to P_i.

Point location: Preprocess an arrangement of n hyperplanes in $\mathbf{R}^d$ so that, given a query point, one can quickly find the face of the arrangement that contains the point. Note that the face need not be d-dimensional. The complexity of a point location algorithm is measured by the query time and the amount of storage needed for the data structure. The time it takes to do the preprocessing is also of importance.

Simplex range searching: Preprocess a set P of n points in $\mathbf{R}^d$ so that, given a query (closed) simplex σ, the size of $P \cap \sigma$ can be quickly evaluated. Simplex range searching refers to a slight generalization of the problem where weights in an additive group or semigroup are assigned to the points and the answer to a query is the sum of all of the weights within σ. This framework allows us to model both the counting and reporting versions of the problem, the latter requiring an explicit enumeration of the points in σ.

CUTTINGS

Clarkson [Cla87] and Haussler and Welzl [HW87] were among the first to introduce the notion of sparsely intersected space partitions for divide and conquer. The definition of an ε-cutting is due to Matoušek [Mat91]. Near-optimal ε-cutting constructions were given found in two dimensions [Aga90, Aga91] and in arbitrary dimension [Mat90, Mat91, Mat95a] (in any dimension). The optimal ε-cutting construction cited below is due to Chazelle. It simplified an earlier design by Chazelle and Friedman to [CF90].

THEOREM 10.3.1 [Cha93]

Given a set H of n hyperplanes in $\mathbf{R}^d$, for any $r > 0$, there exists a $(1/r)$-cutting for H of size $O(r^d)$, which is optimal. The cutting, together with the list of hyperplanes intersecting the interior of each simplex, can be found deterministically in $O(nr^{d-1})$ time.

The standard proof of the theorem is based on a hierarchical construction of independent interest. Roughly, the cutting sought is the last one in a sequence of cuttings $\mathcal{C}_0, \ldots, \mathcal{C}_m$ such that (i) $\mathcal{C}_0$ is of constant size; (ii) for $k > 0$, each simplex of $\mathcal{C}_k$ is enclosed in a unique simplex of $\mathcal{C}_{k-1}$, which itself contains at most a constant number of simplices of $\mathcal{C}_k$; and (iii) for some constant $c > 0$, $\mathcal{C}_k$ is a $(1/c^k)$-cutting of size $O(c^{dk})$.

The simplest application of cuttings is point location in an arrangement of hyperplanes. Consider n hyperplanes in $\mathbf{R}^d$. Given a query point, how fast can we find the cell (or lower-dimensional face) of the arrangement that contains the point? Assuming general position for simplicity, we set $r = n$ in the theorem. From the nesting structure of $\mathcal{C}_0$, $\mathcal{C}_1$, etc, we can locate the query point in $\mathcal{C}_k$ (ie, find the simplex that contains it) in constant time once we know its location within $\mathcal{C}_{k-1}$.

THEOREM 10.3.2 [Cha93]

Point location among n hyperplanes can be done in $O(\log n)$ query time, using $O(n^d)$ preprocessing.

A nice application of cuttings is to the problem of deciding whether there exists any point/line incidence among n lines and n points in the plane. This is often called *Hopcroft's problem.* A well known construction of Erdős provides an arrangement of n lines such that at least n of its vertices are each incident to $\Omega(n^{1/3})$ edges. Choosing these n lines as input to Hopcroft's problem and placing the n points very near the high-degree vertices suggests that to solve the problem should require checking each point against the $\Omega(n^{1/3})$ lines incident to the nearby vertex, for a total of $\Omega(n^{4/3})$ time. This argument can be made rigorous [Eri96]; it offers a strong hint that to beat $\Omega(n^{4/3})$ might not be easy. The bound itself has not been achieved, although an algorithm by Matoušek, based on a subtle use of cuttings, comes near.

THEOREM 10.3.3 [Mat93]

To decide whether n points and n lines in the plane are free of any incidence can be done in time $n^{4/3}\, 2^{O(\log^* n)}$.

SIMPLEX RANGE SEARCHING

Two essential tools in designing data structures for simplex range searching are the *spanning paths* and *simplicial partitions.* We mention the key results about these constructions. As a matter of terminology, we say that a hyperplane cuts a line segment if it intersects it but neither of its endpoints. The n points of a square grid can easily be connected by a path so that no line cuts more than roughly $\sqrt{n}$ edges. The optimal low-cutting spanning path construction of Chazelle and Welzl generalizes this result to any set of points in any dimension.

LEMMA 10.3.4 [CW89]

Any set of n points in $\mathbf{R}^d$ can be ordered as $p_1, \ldots, p_n$, in such a way that no hyperplane cuts more than $cn^{1-1/d}$ segments of the form $p_i p_{i+1}$, for some constant $c > 0$.

Simplicial partitions generalize the notion of a spanning path by considering not just edges, ie, pairs of points, but larger subsets of them. Again, we wish to minimize the cutting number, ie, the number of subsets a hyperplane can "cut through." An optimal construction based on cuttings was discovered by Matoušek.

LEMMA 10.3.5 [Mat92]

Given a set P of n points in $\mathbf{R}^d$ $(d > 1)$, for any integer $1 < r \leq n/2$, there exists a simplicial partition of cutting number $O(r^{1-1/d})$ such that $n/r \leq |P_i| < 2n/r$ for each (P_i, R_i) in the partition.

The partition tree offers a simple solution to simplex range searching. At each node, store the sum of the weights of the points associated with the corresponding region. Given a query simplex σ, we proceed to explore all children v_i of the root and check whether σ intersects the region R_i of v_i: (i) if the answer is yes, but σ does not completely enclose the region R_i of v_i, then we visit v_i and recurse; (ii) if the answer is yes, but σ completely encloses R_i, we simply add to our current weight count the sum of the weights within P_i, which happens to be stored at v_i; (iii) if the answer is no, we do not recurse at v_i.

Applying Lemma 10.3.5 for a large enough constant r yields a partition tree construction which allows us to perform simplex range searching in $O(n^{1-1/d+\varepsilon})$ query time, for any fixed $\varepsilon > 0$, and $O(n)$ storage. A more complex argument by Matoušek gets rid of the ε term in the exponent.

THEOREM 10.3.6 [Mat92]

Given n points in $\mathbf{R}^d$, there exists a linar size data structure with which simplex range searching can be performed in time $O(n^{1-1/d})$ per query.

If superlinear storage is available, then space-time tradeoffs are possible. Chazelle, Sharir, and Welzl [CSW92] proved that simplex range searching with respect to n points in $\mathbf{R}^d$ can be done in $O(n^{1+\varepsilon}/m^{1/d})$ query time, using a data structure of size m, for any $n \leq m \leq n^d$. Matoušek [Mat93] slightly improved the query time to $O(n(\log m/n)^{d+1}/m^{1/d})$, for m/n large enough.

POLYHEDRAL ALGORITHMS

We discuss applications of the discrepancy method to convex hulls, Voronoi diagrams, halfspace intersection, linear programming, and other forms of convex programming.

The problem of computing the convex hull of n points in $\mathbf{R}^d$ reduces by duality to that of computing the intersection of n halfspaces. In addition, to compute the Voronoi diagram (or, equivalently, the Delaunay triangulation) of a finite set of points in Euclidean d-space can be reduced in linear time to a convex hull problem in $(d + 1)$-space. An optimal halfspace intersection algorithm can then be used for both the convex hull and the Voronoi diagram problems. The intersection of n

halfspaces is a convex polyhedron with $O(n^{\lfloor d/2 \rfloor})$ faces (and possibly as many as that).

A simple approach to the halfspace intersection problem is to insert each halfspace one after the other and maintain the current intersection as we go. A simple data structure consisting of a triangulation of the current intersection polyhedron, together with a bipartite graph indicating which hyperplane intersects which cell of the triangulation, is sufficient to make this process efficient. In fact, if the order of insertion is random, then it follows from the work of Clarkson and Shor [CS89] that, with the right supporting data structure, the expected complexity of the algorithm can be made to be optimal. By combining the use of ε-nets, ε-approximations, ε-cuttings, and product set systems, Chazelle [Cha93] showed how to compute the intersection deterministically in optimal time (Theorem 10.3.7); his algorithm was subsequently simplified by Brönnimann, Chazelle, and Matoušek [BCM99].

THEOREM 10.3.7

The polyhedron formed by the intersection of n halfspaces in $\mathbf{R}^d$ can be computed in $O(n \log n + n^{\lfloor d/2 \rfloor})$ time.

As indicated earlier, this result has two important consequences: optimal algorithms for convex hulls and Voronoi diagrams.

THEOREM 10.3.8

The convex hull of a set of n points in $\mathbf{R}^d$ can be computed deterministically in $O(n \log n + n^{\lfloor d/2 \rfloor})$ time.

THEOREM 10.3.9

The Voronoi diagram (or Delaunay triangulation) of a set of n points in $\mathbf{E}^d$ can be computed deterministically in $O(n \log n + n^{\lceil d/2 \rceil})$ time.

Linear programming is the problem of minimizing a linear function $c^T x$, subject to the constraints $Ax \leq b$ and $x \geq 0$, where b is a column vector in $\mathbf{R}^n$ and $c, x \in \mathbf{R}^d$. The discrepancy method can be used to derive a deterministic algorithm for linear programming that is linear in n and singly exponential in d.

The best route to this result is via an abstract formalism, called *LP-type programming*, due to Sharir and Welzl [SW92] (see also [MSW96]) that places the method in a much more general context and allows for even more surprising applications. For example, it can be used to prove that, given n points in, say, $\mathbf{R}^{99}$, the smallest enclosing ellipsoid can be found in $O(n)$ time.

An *LP-type* problem is specified by a pair (H, w), where H is a finite set, whose elements are the "constraints" of the problem, and w is a function mapping certain subsets of H to a totally ordered universe $(W, \leq)$. An element $h \in H$ is said to *violate* a subset $G \subseteq H$ if $w(G) < w(G \cup \{h\})$. A *basis* B of $G \subseteq H$ is a minimal set of constraints with the same cost as G, ie, $w(B) = w(G)$ and $w(C) < w(B)$ for any $C \subset B$. The *combinatorial dimension of* (H, w), denoted by δ, is the maximum size of any basis (of any subset of H). To solve the problem (H, w) is to find a basis of H. We need a few specific assumptions to make computational sense of this framework.

1. Monotonicity. Given any $F \subseteq G \subseteq H$, $w(F) \leq w(G)$.

2. Locality. If $h \in H$ violates $G \subseteq H$, then it violates any basis of G.

3. ORACLE. Given a basis B of some subset of H, let $V(B)$ denote the set of violating constraints. Consider the set system $(H, \mathcal{R})$, where $\mathcal{R}$ is the collection of sets $V(B)$, for all bases B. It is assumed that $(H, \mathcal{R})$ has bounded VC-dimension, and let γ be its shatter function exponent. In practice, γ is either equal or larger than δ. Given any subset $Y \subseteq H$, the oracle computes the set $\mathcal{R}|_Y$ in time $O(|Y|^{\gamma+1})$.

How does linear programming fit into the LP-type framework? For simplicity of explanation, we assume that the optimization function is of the form $(1, 0, \ldots, 0)^T x$, and that the system is feasible: (i) H is the set of n closed halfspaces formed by the inequalities $Ax \leq b$; (ii) $W = \mathbf{R}^d$, ordered lexicographically; (iii) given $G \subseteq H$, $w(G)$ is the unique (lexicographically) minimal point with nonnegative coordinates in the halfspaces of G. A halfspace $h \in H$ violates $G \subseteq H$ if $w(G) < w(G \cup \{h\})$, which means that adding h to G would strictly increase the cost of the optimal solution: Geometrically, the hyperplane corresponding to h cuts off the old solution from the new feasible set. A basis consists of at most d halfspaces, and its combinatorial dimension is d. MONOTONICITY says that throwing in additional constraints cannot improve the optimal solution. LOCALITY means that the violation of a set of constraints can always be witnessed locally by focusing on any one of its bases. The oracle can be implemented easily so as to run in time $O(|Y|^{d+1})$.

Solving an LP-Type Problem

STEP 1. Let $D = \max\{\delta, \gamma\}$. If $|H| \leq cD^4 \log D$, for some suitably large constant c, compute a basis of H by checking all possible j-tuples of constraints in the order $j = 1, \ldots, \delta$, and picking the first one that is not violated by any constraint of H.

STEP 2. Compute a $(1/4D^2)$-net N for $(H, \mathcal{R})$.

STEP 3. Find a basis B of N recursively. Let V be the set of constraints of H that violate B. If $V = \emptyset$, then return B and stop; else add all of the violating constraints to the set N and repeat Step 3.

Assuming that, given any basis B and a constraint $h \in H$, to test whether h violates B or not can be done in time $D^{O(D)}$, which is the case in typical applications, LP-type problems can be solved in time linear in the number of constraints and exponential in the number of variables. Chazelle and Matoušek proved the following:

THEOREM 10.3.10 [CM96]

An LP-type problem (H, w) can be solved in time $|H| \cdot O(D^7 \log D)^D$, where D is the combinatorial dimension or the exponent in the complexity of the oracle, whichever is larger.

We mention two applications of this result, also taken from [CM96]. The first one is a linear deterministic algorithm for linear programming with any fixed number of variables. The second one addresses the complexity of finding the Löwner-John ellipsoid of n points in d-space, ie, the smallest ellipsoid enclosing the n points (which is known to be unique).

THEOREM 10.3.11

Linear programming with n constraints and d variables can be solved in $d^{O(d)}n$ time.

THEOREM 10.3.12

The ellipsoid of minimum volume that encloses a set of n points in $\mathbf{R}^d$ can be computed in time $d^{O(d^2)}n$.

LOWER BOUNDS FOR RANGE SEARCHING

An off-line range searching problem is specified by n points and m regions in $\mathbf{R}^d$. Each point p_i is assigned a *weight* x_i chosen in an additive group or semigroup. The output should be the sum of the weights of the points within each of the m regions. In the on-line version of the problem, the points and weights are preprocessed into a data structure and a query is a region whose weight sum constitutes the output.

From the algebraic perspective of adding weights, off-line range searching can be regarded as the problem of multiplying a fixed matrix by an arbitrary vector. The n points and m ranges form a set system Σ, whose incidence matrix we denote by A. The problem is to compute the map $x \in \mathbf{R}^n \mapsto Ax \in \mathbf{R}^m$. We use a linear circuit model with bounded coefficients. This is a directed acyclic graph whose nodes, the *gates*, have indegree 2. With each gate g is associated two complex numbers α_g, β_g of modulus $O(1)$. The gate g takes two complex numbers a, b as input and outputs $\alpha_g a + \beta_g b$. The size of the circuit is the number of edges. The *complexity* of the matrix A is the size of the smallest circuit for computing $x \mapsto Ax$. We note that the circuit depends only on A and must "work" for any input $x \in \mathbf{R}^n$.

It is not hard to prove that the size of the circuit is $\Omega(\log|\det B|)$, where B is the square submatrix of A whose determinant is largest in absolute value: this is the classical *Morgenstern bound* [Mor73]. A stronger result, due to Chazelle, relates the size of the circuit to the singular values of the matrix A.

SPECTRAL LEMMA [Cha98]

Given an n-by-n $0/1$ matrix A, any circuit for computing Ax has size at least $\Omega(k\log\lambda_k)$, where λ_k is the k-th largest eigenvalue of A^TA.

A slightly weaker formulation of the spectral bound was given by Chazelle and Lvov [CL00]. It involves only the traces of A^TA and its square. This has the huge advantage that every component of the formula as a simple combinatorial interpretation: the trace of A^TA (of its square) counts the number of ones (resp. rectangles of ones) in A.

TRACE LEMMA [CL00]

Given an n-by-n $0/1$ matrix A, any circuit for computing Ax has size

$$\Omega_\varepsilon\Big(n\log\Big(\mathrm{tr}\,M/n - \varepsilon\sqrt{\mathrm{tr}\,M^2/n}\Big)\Big),$$

where $M = A^TA$ and $\varepsilon > 0$ is an arbitrarily small constant.

The bounds can be made more general to accommodate a few "help" gates, ie, gates that can compute any function whatsoever [Cha98]. The spectral and trace lemmas have been used to derive bounds for a number of classical range searching problems. The monotone model of computation, where essentially subtractions are disallowed, has also been investigated. We mention the main results below and explain their meaning. The proofs rely heavily on tools from discrepancy theory; in particular, constructions of low-discrepancy point sets and techniques from harmonic analysis to analyze the spectrum of geometric incidence matrices.

TABLE 10.3.1 Circuit lower bounds for range searching.

TYPE	GENERAL	MONOTONE
axis parallel boxes	$\Omega(n \log\log n)$	$\Omega(n(\log n/\log\log n)^{d-1})$
simplices	$\Omega(n \log n)$	$\widetilde{\Omega}(n^{2-2/(d+1)})$
lines	$\Omega(n \log n)$	$\Omega(n^{4/3})$

The table indicates some of the lower bounds known to date. In all cases, the problem consists of n points in $\mathbf{R}^d$ and n regions whose type is indicated in the first column. The GENERAL column refers to the circuit model discussed above. The proofs were given for $\alpha_g, \beta_g \in \{-1, 0, 1\}$ but extend trivially to any complex numbers with bounded modulus. The bounds for axis-parallel boxes [Cha97] and simplices [Cha98] were proven in dimension 2 and, hence, in any higher dimension. In the case of axis-parallel boxes, the lower bound jumps to $\Omega(n \log n)$ in dimension $\Theta(\log n)$ [CL01]. The bound for lines was proven in [CL00].

The MONOTONE column assumes that $\alpha_g, \beta_g \in \{0, 1\}$ at each gate of the circuit. The notation $\widetilde{\Omega}(f(n))$ refers to $\Omega(f(n)/(\log n)^{O(1)})$. The bounds for axis-parallel boxes and simplices were given in [Cha97]. The bound for lines is mentioned in [Cha00]. All three bounds are essentially optimal in that model.

It is a fascinating open question how the wide gap between general and nonmonotone complexity is to be resolved. For example, is line range searching in $O(n \log n)$ or $O(n^{4/3})$? Most of the lower bounds for the monotone case have nearly matching upper bounds; in other words, to make effective use of nonmonotone computation seems very difficult. This has led to the widely held belief that the monotone bounds are the true answers. Recent work by Chazelle [Cha01] casts doubt on this conjecture. Using grid points and line queries that bounce off the grid boundary, the general complexity of the problem is shown to be $\Theta(n \log n)$, while the monotone complexity is $\Theta(n^{3/2})$.

In the on-line version of range searching, the n points are preprocessed so that, given a query region, the sum of the weights of the points in the region can be computed quickly. The following bounds, established by Chazelle in the monotone model, are essentially optimal. Both of them make heavy use of low-discrepancy constructions for point sets in bounded-dimensional space, as well as related constructions arising from Heilbronn's problem [Rot51].

THEOREM 10.3.13 [Cha89]

Given n points in $\mathbf{R}^d$, on-line simplex range searching requires $\widetilde{\Omega}(n/m^{1/d})$ query time, using a data structure of size m.

THEOREM 10.3.14 [Cha90]

Given n points in $\mathbf{R}^d$, on-line range searching with axis-parallel box queries requires $\Omega(\log n/\log(2m/n))^{d-1}$ per query, using a data structure of size m.

FURTHER READING

Many aspects of the discrepancy method, including nongeometric ones, are covered in [Cha00]. The related topic of derandomization is surveyed in [Mat96]. The main texts on discrepancy theory are [BC87, Nie92, DT97, Mat99]; see also the survey [BS95].

RELATED CHAPTERS

Chapter 11: Geometric discrepancy theory and uniform distribution
Chapter 21: Convex hull computations
Chapter 22: Voronoi diagrams and Delaunay triangulations
Chapter 35: Range searching
Chapter 38: Randomized algorithms
Chapter 39: Derandomization
Chapter 44: Linear programming in low dimensions

REFERENCES

[Aga90] Agarwal, P.K. Partitioning arrangements of lines II: Applications, *Disc. Comput. Geom.* 5 (1990), 533–573.

[Aga91] Agarwal, P.K. Geometric partitioning and its applications, in *Computational Geometry: Papers from the DIMACS Special Year,* eds., Goodman, J.E., Pollack, R., Steiger, W., Amer. Math. Soc., 1991.

[ARS99] Alon, N., Rónyai, L., Szabó, L. Norm-graphs: variations and applications, *J. Combinatorial Theory B* 76 (1999), 280–290.

[Ass83] Assouad, P. Densité et dimension, *Annales de l'Institut Fourier* 33 (1983), 233–282.

[BC87] Beck, J., Chen, W.W.L. *Irregularities of Distribution*, Cambridge Tracts in Mathematics, 89, Cambridge University Press, 1987.

[BS95] Beck, J., Sós, V.T. Discrepancy theory, in *Handbook of Combinatorics,* Chap. 26, eds., Graham, R.L., Grötschel, M., Lovász, L., North-Holland, 1995, pp. 1405–1446.

[BCM99] Brönnimann, H., Chazelle, B., Matoušek, J. Product range spaces, sensitive sampling, and derandomization, *SIAM J. Comput.* 28 (1999), 1552–1575.

[Cha89] Chazelle, B. Lower bounds on the complexity of polytope range searching, *J. Amer. Math. Soc.* 2 (1989), 637–666.

[Cha90] Chazelle, B. Lower bounds for orthogonal range searching: II. The arithmetic model, *J. ACM* 37 (1990), 439–463.

[Cha93] Chazelle, B. Cutting hyperplanes for divide-and-conquer, *Disc. Comput. Geom.* 9 (1993), 145–158.

[Cha93] Chazelle, B. An optimal convex hull algorithm in any fixed dimension, *Disc. Comput. Geom.* 10 (1993), 377–409.

[Cha97] Chazelle, B. Lower bounds for off-line range searching, *Disc. Comput. Geom.* 17 (1997), 53–65.

[Cha98] Chazelle, B. A spectral approach to lower bounds with applications to geometric searching, *SIAM J. Comput.* 27 (1998), 545–556.

[Cha00] Chazelle, B. *The Discrepancy Method: Randomness and Complexity*, Cambridge University Press, hardcover 2000, paperback 2001.

[Cha01] Chazelle, B. The power of nonmonotonicity in geometric searching, *Proc. 18th Annual ACM Symp. Comput. Geom.* (2002), 88–93. To appear in Disc. Comput. Geom.

[CF90] Chazelle, B., Friedman, J. A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.

[CL00] Chazelle, B., Lvov, A. A trace bound for the hereditary discrepancy, *Disc. Comput. Geom.* 26 (2001), 221–231.

[CL01] Chazelle, B., Lvov, A. The discrepancy of boxes in higher dimension, *Disc. Comput. Geom.* 25 (2001), 519–524.

[CM96] Chazelle, B., Matoušek, J. On linear-time deterministic algorithms for optimization problems in fixed dimension, *J. Algorithms* 21 (1996), 579–597.

[CSW92] Chazelle, B., Sharir, M., Welzl, E. Quasi-optimal upper bounds for simplex range searching and new zone theorems, *Algorithmica* 8 (1992), 407–429.

[CW89] Chazelle, B., Welzl, E. Quasi-optimal range searching in spaces of finite VC-dimension, *Disc. Comput. Geom.* 4 (1989), 467–489.

[Cla87] Clarkson, K.L. New applications of random sampling in computational geometry, *Disc. Comput. Geom.* 2 (1987), 195–222.

[CS89] Clarkson, K.L., Shor, P.W. Applications of random sampling in computational geometry, II, *Disc. Comput. Geom.* 4 (1989), 387–421.

[DT97] Drmota, M.,Tichy, R.F., *Sequences, Discrepancies and Applications*, Lecture Notes in Mathematics, Vol. 1651, Springer, 1997.

[Eri96] Erickson, J. New lower bounds for Hopcroft's problem, *Disc. Comput. Geom.* 16 (1996), 389–418.

[Fre81] Fredman, M.L. Lower bounds on the complexity of some optimal data structures, *SIAM J. Comput.* 10 (1981), 1–10.

[HW87] Haussler, D., Welzl, E. ε-nets and simplex range queries, *Disc. Comput. Geom.* 2 (1987), 127–151.

[Joh74] Johnson, D.S. Approximation algorithms for combinatorial problems, *J. Comput. Syst. Sci.* 9 (1974), 256–278.

[KPW92] Komlós, J., Pach, J., Woeginger, G. Almost tight bounds for ε-nets, *Disc. Comput. Geom.* 7 (1992), 163–173.

[Lov75] Lovász, L. On the ratio of optimal integral and fractional covers, *Disc. Math.* 13 (1975), 383–390.

[Mat90] Matoušek, J. Construction of ε-nets, *Disc. Comput. Geom.* 5 (1990), 427–448.

[Mat91] Matoušek, J. Cutting hyperplane arrangements, *Disc. Comput. Geom.* 6 (1991), 385–406.

[Mat92] Matoušek, J. Efficient partition trees, *Disc. Comput. Geom.* 8 (1992), 315–334.

[Mat93] Matoušek, J. Range searching with efficient hierarchical cuttings, *Disc. Comput. Geom.* 10 (1993), 157–182.

[Mat95a] Matoušek, J. Approximations and optimal geometric divide-and-conquer, *J. Comput. Syst. Sci.* 50 (1995), 203–208.

[Mat95b] Matoušek, J. Tight upper bounds for the discrepancy of halfspaces, *Disc. Comput. Geom.* 13 (1995), 593–601.

[Mat96] Matoušek, J. Derandomization in computational geometry, *J. Algorithms* 20 (1996), 545–580.

[Mat97] Matoušek, J. On discrepancy bounds via dual shatter functions, *Mathematika* 44 (1997), 42–49

[Mat99] Matoušek, J. *Geometric Discrepancy: An Illustrated Guide*, Algorithms and Combinatorics, 18, Springer, 1999.

[MSW96] Matoušek, J., Sharir, M., Welzl, E. A subexponential bound for linear programming, *Algorithmica* 16 (1996), 498–516.

[MWW93] Matoušek, J., Welzl, E., Wernisch, L. Discrepancy and ε-approximations for bounded VC-dimension, *Combinatorica* 13 (1993), 455–466.

[Mor73] Morgenstern, J. Note on a lower bound of the linear complexity of the fast Fourier transform, *J. ACM* 20 (1973), 305–306.

[Nie92] Niederreiter, H. *Random Number Generation and Quasi-Monte Carlo Methods*, CBMS-NSF, SIAM, Philadelphia, PA, 1992.

[PA95] Pach, J., Agarwal, P.K. *Combinatorial Geometry*, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., 1995.

[Rot51] Roth, K.F. On a problem of Heilbronn, *J. London Math. Soc.* 26 (1951), 198–204.

[Sau72] Sauer, N. On the density of families of sets, *J. Combinatorial Theory A* 13 (1972), 145–147.

[SW92] Sharir, M., Welzl, E. A combinatorial bound for linear programming and related problems, *Proc. 9th Annual Symp. Theoret. Aspects Comput. Sci.,* LNCS, 577, Springer-Verlag, 1992, pp. 569–579.

[She72] Shelah, S. A combinatorial problem; stability and order for models and theories in infinitary languages, *Pac. J. Math.* 41 (1972), 247–261.

[VC71] Vapnik, V.N., Chervonenkis, A.Ya. On the uniform convergence of relative frequencies of events to their probabilities, *Theory of Probability and its Applications* 16 (1971), 264–280.