



HAL
open science

Approximations of Points: Combinatorics and Algorithms

Nabil Mustafa

► **To cite this version:**

Nabil Mustafa. Approximations of Points: Combinatorics and Algorithms. Computational Geometry [cs.CG]. Université Paris-Est, 2013. tel-01062825

HAL Id: tel-01062825

<https://theses.hal.science/tel-01062825v1>

Submitted on 10 Sep 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS-EST

Approximations of Points: Combinatorics and Algorithms

Nabil Mustafa

Mémoire d'habilitation à diriger des recherches

(Spécialité Informatique)

M. Victor CHEPOI	Professeur, Université de la Méditerranée, Marseille	Rapporteur
M. Kurt MEHLHORN	Professeur, MPI für Informatik & Saarland University	Rapporteur
M. Emo WELZL	Professeur, ETH Zurich	Rapporteur
M. Gilles BERTRAND	Professeur, ESIEE Paris	Examineur
Mme. Claire MATHIEU	Directeur de Recherche, ENS Paris	Examineur
M. Dominique PERRIN	Directeur Général, ESIEE Paris	Directeur d'habilitation
M. Michel POCCHIOLA	Professeur, Université Pierre et Marie Curie	Examineur
M. Gilles SCHAEFFER	Directeur de Recherche, École Polytechnique	Examineur

November 2, 2013

This page was intentionally left blank

Foreword

My steadfast interest has been in *geometry*, exploring the interaction of the structure of geometric shapes with computation. To that end, I have spent the past decade learning ideas and using techniques from several areas: algorithms, computational geometry, GPU computing, computational biology, computer graphics, databases, visualization, algebraic topology, discrete geometry and randomized analysis.

Given available space (and reader patience!), my habilitation thesis focuses on my contributions in a particular area, and so gives only a partial account of my research interests and activities. This foreword fills this gap by giving a broader overview of my research activities for the past several years.

* * *

My work for the past fifteen years has been, broadly, in the area of geometric computing, i.e., aspects of computation on geometric data. Geometric data is ubiquitous – from Global Positioning Systems (GPS) to Medical Imaging to Computational Biology to Geographical Information Systems to 3D scanning. Successful understanding and processing of such data involves the interface of several key areas, requiring expertise in Mathematics, Algorithms, Geometry and Graphics & Visualization. Computational geometry, an amalgamation of the above areas, branched out as a separate field in the 1970s, with the aim of providing a rigorous mathematical and algorithmic basis for the study and computation of geometric data.

One of the fundamental themes in discrete and computational geometry is the notion of *approximation* of geometric data. There are, broadly, two related ways to approximate geometry: either replace the data with a smaller but similar data (e.g., its median, centroid) which ‘combinatorially’ approximates the original set P . Or compute some descriptive quantity (e.g., diameter, width, variance) over the entire set P .

I now turn to each in more detail. A complete list of my publications can be found on page 83 .

When the geometric data consists of a set P of n points in \mathbb{R}^d , the goal of data-depth measures is to generalize the idea of the median in higher dimensions; over the past 50 years, several such measures have been proposed and studied. Part of my work has been a study of their combinatorial and algorithmic aspects – designing efficient algorithms for the computation of *Tukey depth* [15], resolution of the Oja-depth depth conjecture in \mathbb{R}^2 and improving bounds for general dimensions [33], improved algorithms for Ray-Shooting depth [31], improved combinatorial bounds for the simplicial-depth problem [11], introducing and developing the notion of center-disks [10], proposing graphics hardware based algorithms for several of the above depth-measures (simplicial-depth, oja-depth, colorful variants of these depth measures) [16], and finally, proposing (and partially proving) a general conjecture that unifies all of these above measures under the *Centerflat Conjectures* [32].

Similarly, there are several notions of approximating P with more than one point. The most well-known and useful one is that of ε -nets, where one would like to compute a small-sized subset of P

that is a hitting set for all large-enough subsets of P induced by some geometric object (for example, under balls, rectangles). One prominent case is the so-called *weak ε -net problem*, a long-standing open problem on approximating convex sets. I proposed optimal constructions for small-sized weak ε -nets [27]. For the general problem, I introduced the use of random-sampling for computing weak ε -nets [23] which showed a connection between ε -nets and weak ε -nets. This work required the generalization of some basic theorems in discrete geometry, which was presented in [30]. Finally, some near-optimal generalizations of the classical ε -net theorem were presented in [24].

Apart from the combinatorial bounds on the size of the hitting sets, I have also worked on algorithms for computing hitting-sets, and related problems. An open problem of finding a polynomial-time approximation scheme for hitting-sets for disks in the plane was resolved in [28, 29]; this paper showed the use of provably good local-search for combinatorial geometric problems via locality-graphs, an idea that was used by several later papers. I also presented approximation algorithms for computing maximum independent sets in intersection graphs of geometric objects in the plane [4, 6]. Other works include streaming algorithms for computation of minimum-width, smallest enclosing ball and other problems [3], and the first constant-factor approximation algorithm for the well-known group TSP problem [13].

When the data is ordered, e.g., a polygonal curve, I presented approximation algorithms for their simplification under Hausdorff and Frechet distance metrics [1, 2]; here the motivating application was simplifying protein backbones for structural similarity. This led to the harder problem of simplifying two-dimensional cartographic maps, and in further work [19, 20], we proposed real-time topology-preserving algorithms for dynamic simplification and visualization of maps. This work forms the theoretical basis for a patent (US Patent Nr. 10/021,645).

Geometric data comes from many different sources, so part of my published work spans these areas: Databases, where algorithms for projective clustering of high-dimensional point-set data were proposed [5]; Computational Biology, where in [7] we resolved several of the problems related to comparing protein backbones via the so-called *contact-map overlap measure*; Computational Chemistry, where we presented the construction and main ideas that were then used by others to resolve the Weiner-index conjecture [9, 8]; Networks, where the frequency assignment problem for wireless stations was studied [14], as well as the majority influence problem [21, 22].

Our work in [3, 15, 16, 19, 20] was one of the earliest works exploiting the graphics hardware for geometric computation, and demonstrated its untapped potential; by now, graphics computing is a large mainstream area. Our work was presented at SIGGRAPH course notes in 2002, entitled “Interactive Geometric Computations with Graphics Hardware”.

* * *

This thesis presents a subset of the above work, that related to combinatorial and algorithmic aspects of approximations of point-set data in \mathbb{R}^d . In particular, this exposition relies on the material from the following papers. I should note that besides Chapter 5 which includes joint work with my PhD advisor, the co-authors of all the remaining papers are young researchers (Hans Raj Tiwary, Rajiv Raman) or PhD students (or were when we started collaborating; Saurabh Ray, Abdul Basit, Mudassir Shabbir, Sarfraz Raza, Daniel Werner).

- *Geometric separators*, recent unpublished work with Rajiv Raman and Saurabh Ray (Chapter 2),
- *The general weak ε -net problem* [26, 23], with Saurabh Ray (Chapter 3),
- *Near-optimal generalization of Caratheodory's Theorem*, with Saurabh Ray (Chapter 4),
- *Independent sets in Intersection-graphs* [4, 6], with Pankaj Agarwal (Chapter 5),
- *Hitting-sets for disks in the plane* [28, 29], with Saurabh Ray (Chapter 6),
- *Optimal bounds for small weak ε -nets* [25, 27], with Saurabh Ray (Chapter 7),
- *Improved bounds for Simplicial Depth* [11, 12], with Abdul Basit, Saurabh Ray and Sarfraz Raza (Chapter 8),
- *Optimal bounds for Oja-depth* [33], with Hans Raj Tiwary and Daniel Werner (Chapter 9),
- *Unifying data-depths via the Centerflat Conjectures* [32], with Saurabh Ray and Mudassir Shabbir (Chapter 1).

This page was intentionally left blank

Abstract

At the core of successful manipulation and computation over large geometric data is the notion of *approximation*, both *structural* and *computational*. The focus of this thesis will be on the combinatorial and algorithmic aspects of approximations of point-set data P of n points in \mathbb{R}^d .

Part I of this thesis considers the problem of approximating P with a small-sized subset. There we concentrate on the long-standing open problem of weak ε -nets: given a set P of n points in \mathbb{R}^d and a parameter $\varepsilon > 0$, one would like to compute a small-sized subset $Q \subset \mathbb{R}^d$ such that any convex set containing at least εn points of P contains at least one point of Q ; i.e., Q is a hitting-set for sets of size at least εn induced by convex objects in \mathbb{R}^d . Improving earlier work, we first present optimal constructions of Q when ε is large; in particular, we present optimal bounds when Q consists of two points, generalizing the classical centerpoint theorem. For the general weak ε -net problem, the current-best bounds on the size of Q are $O(1/\varepsilon^d)$. We will show that picking $O(1/\varepsilon \log 1/\varepsilon)$ points R randomly from P give sufficient information to be able to construct a weak ε -net solely from R . Along the way, this requires re-visiting and generalizing some basic theorems of convex geometry, such as the Carathéodory's theorem.

Part II of this thesis is a consideration of the algorithmic aspects of these problems. We first present a PTAS for computing hitting-sets for disks in the plane, thus closing a 30-year old open problem. Of separate interest is the technique, an analysis of local-search via locality graphs. A further application of this technique is then presented in computing independent sets in intersection graphs of rectangles in the plane.

Part III of this thesis deals with notions of geometric data depth, in particular those of geometric medians in \mathbb{R}^d . Given P , the goal is to compute a point $q \in \mathbb{R}^d$ that is at the 'combinatorial center' of P ; i.e., a natural analog of the concept of median in higher dimensions. Over the past 50 years several such measures have been proposed, and we will re-examine some of the central ones: Tukey depth, Simplicial depth, Oja depth and Ray-Shooting depth. For each of these measures, we improve the state-of-the-art by presenting either faster algorithms or improved combinatorial bounds. Finally we give a geometric framework that unifies these previously considered separates measures.

This page was intentionally left blank

Contents

1	A Detailed Overview	1
Part I Around ϵ-nets		
2	Geometric Separators	21
3	The General Weak ϵ -net problem	27
4	A Generalization of Carathéodory's Theorem	35
Part II Algorithms		
5	Independent sets in Intersection Graphs of Rectangles	43
6	A PTAS for the Geometric Hitting-Set Problem for Disks	49
Part III Data Depth		
7	A Generalization of the Centerpoint Theorem	61
8	Simplicial Depth	67
9	Oja Depth	73
Appendices		
A	List of Basic Theorems	81

This page was intentionally left blank

1. A Detailed Overview

Mathematics is amazingly compressible: you may struggle a long time, step by step, to work through some process or idea from several approaches. But once you really understand it and have the mental perspective to see it as a whole, there is a tremendous mental compression. You can file it away, recall it quickly and completely when you need it, and use it as just one step in some other mental process. The insight that goes with this compression is one of the real joys of mathematics.

William Thurston

THIS chapter, the longest one, gives a detailed account of the results in this thesis: background and context of the problems, historical progression of ideas, relations between seemingly different results, a unifying view of some of the problems, significance of the results and their algorithmic implications. This will be accomplished keeping three other aims in mind: to provide intuition behind the various technical statements, to use this opportunity to more deeply explain fundamental techniques, and to present several conjectures and plausible directions of further research. We hope that this will give this thesis usefulness beyond just a presentation of a specific set of results.

Towards these ends, some material from various chapters has been moved to this one. The remaining chapters are then direct technical proofs of the statements claimed here. While reading here, the reader can visit the appropriate chapters for the technical proofs. A list of well-known theorems mentioned or used can be reviewed in Appendix A.

★ ★ ★

ε -nets. One of the fundamental ways to combinatorially approximate a set P of points in \mathbb{R}^d is to capture the desired properties (depending on the application) of P with a smaller-sized subset of \mathbb{R}^d ; e.g., when P is a set of points in \mathbb{R} , the notion of *mean* and *median* are two natural widely-used measures to capture the spread of the data. The more geometric of these measures, the mean, generalizes in a straightforward way to higher dimensions as the *centroid*; the appropriate generalization of the combinatorial one, the median, is less clear. For now we mention one fundamental generalization, the 1947 *centerpoint theorem* of Rado [Rad47]¹:

¹The third part of this thesis deals with various notions of combinatorial approximations of P with a single point.

Centerpoint Theorem. *Given any set P of n points in \mathbb{R}^d , there exists a point $q \in \mathbb{R}^d$, not necessarily in P , such that any closed half-space containing q contains at least $n/(d+1)$ points of P . Any such point q is called a centerpoint of P .*

One proof of this statement follows from the realization that it is equivalent to the fact that all convex objects containing greater than $dn/(d+1)$ points of P can be hit by one point². Helly's theorem then implies the centerpoint theorem.

This equivalence paves the way for the natural extension of approximating with k points instead of just one: given an integer k , what is the smallest $\varepsilon > 0$ such that given any set P of n points in \mathbb{R}^d , there exist a set Q of k points such that any convex object containing at least εn points of P is hit by Q . Equivalently, one can fix ε and ask for the smallest-sized set $Q \subset \mathbb{R}^d$ so that any convex object containing at least εn points of P contains at least one point of Q . This is the well-known weak ε -net problem.

More generally, one can ask questions regarding small-sized hitting sets for subsets induced by other kinds of geometric objects. Broadly, given a set system (X, \mathcal{F}) where X is a base set and \mathcal{F} is a family of subsets of X , the ε -net problem asks for a small-sized subset $X' \subset X$ such that for every set $S \in \mathcal{F}$ containing at least $\varepsilon|X|$ elements, $X' \cap S \neq \emptyset$. Examples of geometric set-systems include when X is a set of points in \mathbb{R}^d , and \mathcal{F} is defined by containment by geometric objects such as halfspaces or balls or rectangles. Note here that unlike the convex case, we require the hitting-set to be a subset of the base set.

The main question (with many implications), first proposed in Haussler-Welzl [HW87], concerns bounding the sizes of ε -nets. Their remarkable and celebrated result showed that if the set-system satisfied a certain combinatorial condition (VC-dimension is at most d), then picking a random sample of size $O(d/\varepsilon \log d/\varepsilon)^3$ yields an ε -net with high probability, the size being *independent* of $|X|$! The condition of having finite VC-dimension is satisfied by set systems induced by many geometric objects: disks, half-spaces, k -sided polytopes, r -admissible set of regions etc. in \mathbb{R}^d , and in general, geometric objects defined by bounded polynomial constraints. Furthermore, this bound is tight if one only relies on VC-dimension [KPW92]; consequently VC-dimension theory is inadequate for proving asymptotically better upper-bounds for geometric set-systems. For example, given a set of points in \mathbb{R}^d , set systems defined by containment by balls, halfspaces, rectangles have been of fundamental combinatorial and computational interest. So over the past two decades, a number of specialized techniques have been developed for such set systems for a number of geometric settings [PW90, MSW90, PR08, CF90, CV07, AES10, Var10, CGKS12].

In a recent breakthrough, Alon [Alo12] proved a super-linear lower-bound on the sizes of ε -nets for set-systems induced by lines in the plane. Alon's observation was that any Ramsey-theoretic statement immediately implies a lower-bound on ε -nets for a related set-system. For an illustration (Alon uses a different Ramsey-type statement), consider the consequence of the following Ramsey-type statement on arithmetic progressions: given an integer k , there exists an integer N_k (function of k) such that any subset of $A = \{1, \dots, N_k\}$ of size $N_k/2$ contains an arithmetic progression of size k . Now suppose we want to hit all arithmetic progressions of size k in A ; i.e., an ε -net S

²Clearly a centerpoint q must hit all convex sets containing greater than $dn/(d+1)$ points, for otherwise the halfspace containing q and separating q from any such unhit convex-set contains less than $n/(d+1)$ points. A similar argument shows that any point q hitting all such convex-sets must be a centerpoint.

³This bound was later improved in [KPW92] to a near-optimal bound of $(1 + o(1))(\frac{d}{\varepsilon} \log(1/\varepsilon))$.

for $\varepsilon = k/N_k$. Now the statement above implies that S cannot have size linear in $O(1/\varepsilon)$ (i.e., $O(N_k/k)$): S must have size at least $N_k/2$, as otherwise the set $A \setminus S$ would contain an arithmetic progression of length k that is not hit by S . This gives a non-linear lower-bound depending on the exact function relating k and N_k ⁴. With a quantitatively stronger Ramsey-type statement, Pach-Tardos [PT11] showed that any ε -net for halfspaces must have size $\Omega(1/\varepsilon \log 1/\varepsilon)$ in dimension 4 and higher, thus closing a long-standing open problem.

One reason for the usefulness of ε -nets is that they capture precisely and provably properties one would want from a random sample (a ‘well spread out’ subset). So there have been numerous applications of ε -nets, both combinatorial (the spanning-tree theorem, partition theorem, use in geometric discrepancy) and algorithmic (divide-and-conquer, hitting-set algorithms as we will see in the second part).

We give an application of ε -nets to *geometric separators* that illustrates both their combinatorial and algorithmic usefulness. Separators are substructures whose removal partitions the original structure into smaller non-interfering substructures, thus allowing independent computation of smaller subproblems. A classical example is the planar graph separator theorem [LT80]: any planar graph G has a small-sized subset that partitions the graph into two roughly equal-sized independent pieces. Specifically, given a planar graph $G = (V, E)$, there exists a subset $S \subset V$ of size $O(\sqrt{n})$ such that $V \setminus S$ can be partitioned into two sets V_1, V_2 with no edge between vertices in V_1 and V_2 in E .⁵

Very recently a new separator theorem was proven and used for the problem of approximating the maximum independent set in the intersection graphs of geometric objects in the plane⁶. Adamaszek-Wiese [AW13, AW14] showed the following theorem: given a set S of n weighted disjoint line segments in the plane of total weight W , and a parameter $\delta > 0$, there exists a piecewise linear simple closed curve \mathcal{C} with $O((1/\delta)^6)$ vertices such that *i*) the weight of the line segments in S intersecting \mathcal{C} is at most δW , and *ii*) the weight of the line segments completely inside or outside \mathcal{C} is at most $2W/3$. This theorem then directly implies, via dynamic programming, a quasi-polynomial time approximation scheme (QPTAS) for computing independent-sets in the intersection graphs of geometric objects (defined by at most polylogarithmically many vertices) in the plane.

Their separator proof is long, and technically messy, using a theorem of Arora *et al.* [AGK⁺98]. As an illustration of the versatility of ε -nets, we note that an improved optimal separator theorem follows easily from ε -nets and a variant of the planar graph separator theorem stated earlier:

Theorem (Chapter 2). *Given a set S of n disjoint weighted line segments in the plane (with total weight W) and a parameter $\delta > 0$, there exists a piecewise linear simple closed curve \mathcal{C} with $O(1/\delta)$ vertices such that *i*) the total weight of segments intersecting \mathcal{C} is at most δW , and *ii*) the total weight of the line segments completely inside or outside \mathcal{C} is at most $2W/3$. Furthermore this is optimal, in the sense that any curve \mathcal{C} satisfying these two properties must have size $\Omega(1/\delta)$.*

Besides improving the running time of the algorithms of [AW13, AW14] by polynomial factors, the

⁴This set-system can be viewed as geometric by placing A as points in \mathbb{R}^2 ($p \in \mathbb{R} \rightarrow (p, 0) \in \mathbb{R}^2$); then the arithmetic progressions are induced by sinusoids in the plane.

⁵Later we will see the use of this separator theorem for the resolution of the problem of a polynomial-time approximation scheme (PTAS) for the hitting-set problem.

⁶Given a set $\{O_1, \dots, O_n\}$ of n objects in some Euclidean space, the intersection graph has a vertex for each object, and an edge between two vertices iff the two corresponding objects intersect.

above statement can be generalized to more general objects like x -monotone curves in the plane, intersecting curves etc..

* * *

Weak ε -nets. Returning to the convex case, given a set P of n points in \mathbb{R}^d , and a parameter $\varepsilon > 0$, call a set $Q \subset \mathbb{R}^d$ a *weak ε -net* for P if any convex object containing at least εn points of P is hit by Q . The previous discussion enables us to view weak ε -nets from two perspectives: as a generalization of the centerpoint theorem to subsets, and as ε -nets where the set-system is induced by convex objects.

Getting optimal bounds on the sizes of weak ε -nets has been a long-standing open problem⁷, as the ε -net theorem of Haussler-Welzl does not apply (the set systems induced by convex objects have unbounded VC-dimension). Alon *et al.* [ABFK92] showed that one can derive weak ε -nets from the iterative use of centerpoints. This gave the first construction of weak ε -nets of size *independent* of n , with a bound of $O(1/\varepsilon^{d+1})$. This was subsequently improved by Chazelle *et al.* [CEG⁺93] to $O(1/\varepsilon^d \text{polylog}(1/\varepsilon))$; more than a decade later, an elegant new proof of this statement was given by Matousek-Wagner [MW04]. On the other hand, Bukh *et al.* [BMN09] proved a lower bound of $\Omega(1/\varepsilon \log^{d-1} 1/\varepsilon)$. Quoting Matousek-Wagner [MW04]: “There seems to be no convincing reason why $f(d, \varepsilon)$ [size of the weak ε -net in \mathbb{R}^d] should be substantially super-linear in $1/\varepsilon$ ”. As we will see, for a variety of reasons, a first aim may be the middle-ground of $O(1/\varepsilon^{\lfloor d/2 \rfloor})$.

To understand the fundamental reason why progress has been lacking, and there is a wide gap between the upper- and lower- bounds, consider the proof of the ε -net theorem: each set to be hit contains a large fraction of the points (at least εn), and so a random sample has a high probability of being a hitting-set *as long as the number of sets is small enough for the union-bound to work*. This idea works when sets are defined by half-spaces, or balls or any algebraically bounded object, or have finite VC-dimension. For all these, the number of sets are polynomial in $|X|$, and so trivially a random sample of size $O(1/\varepsilon \log |X|)$ is a hitting set (the subtle analysis of Haussler-Welzl is able to achieve the bound of $O(1/\varepsilon \log 1/\varepsilon)$ by essentially using the ‘hereditary’ property of the VC-dimension of a set-system). This sampling approach seems doomed for convex objects as the number of induced subsets are exponentially large (then one requires the probability of not hitting a set to be exponentially small as well).

We propose a new program towards the resolution of this problem: we will show that given the set P of n points in \mathbb{R}^d , it is possible to pick a subset of only $O(1/\varepsilon \log 1/\varepsilon)$ points of P from which a weak ε -net can be constructed (though the final size of the weak ε -net is still exponential in d). The high-level idea is the following: pick a random sample R of size $O(1/\varepsilon \log 1/\varepsilon)$ from P . Now suppose a convex object C containing εn points of P is not hit by this sample. Then it must be that no set of k halfspaces separates⁸ C from points of R ; otherwise the intersection of these halfspaces contains C and so contains at least εn points without containing any point of R , an unlikely event

⁷We refer the reader to a good history and review of the problem by Gil Kalai: <http://terrytao.wordpress.com/2007/04/22/gil-kalai-the-weak-epsilon-net-problem/>.

⁸We say a point p can be *separated* from a convex set C if there exists a hyperplane h with C and p in the interior of the two different halfspaces defined by h .

as intersection of k halfspaces has finite VC-dimension. Intuitively, this means that some subset R_C of R is ‘close’ to the boundary of C : specifically, it will be shown that there exists a large-enough set $R_C \subset R$ such that the convex-hull of every $(d + 1)$ tuple of R_C intersects C . We then show that there exists an appropriate set of ‘product’ points Q over R that hit every convex object C containing at least εn points of P .

Theorem (Chapter 3). *Given a set P of n points in \mathbb{R}^d , and a parameter $\varepsilon > 0$, there exists a set R of $O(1/\varepsilon \log 1/\varepsilon)$ points of P such that a weak ε -net Q for P w.r.t. convex sets can be constructed from points in R .*

The proof establishes the use of random-sampling for computing weak ε -nets: basically using centerpoints over a random sample of size $O(1/\varepsilon \log 1/\varepsilon)$ gives, with high probability, a weak ε -net! In fact, the proof also establishes a connection between weak ε -nets and ε -nets, showing that weak nets can be constructed from ε -nets.

We use two ‘product’ functions over points of R to construct Q : Radon points, and centerpoints. Given P , we say a point in \mathbb{R}^d has *descriptive complexity* k if it is completely determined by at most k points of P . For example, a Radon point of a $(d + 2)$ -tuple in P trivially has descriptive complexity $d + 2$. It can also be shown that there exists a centerpoint of P with d^2 descriptive complexity⁹. We ask the following question: does there exist an *approximate* centerpoint (in the sense that any closed halfspace containing such a point contains at least n/c_d points of P , where c_d can be any function of d) of descriptive complexity less than d points¹⁰? Our proof yields the following connection: if the descriptive complexity of an approximate centerpoint is t , then one can construct weak ε -nets of size $O(1/\varepsilon^t \log^t 1/\varepsilon)$. Thus our approach directly relates the size of the weak ε -nets to the descriptive complexity of basic product functions.

One of the bottlenecks in the above program is the weak property of ‘closeness’ satisfied by R_C , i.e., that the convex-hull of every $d + 1$ subset of R_C intersects C . This blocks the approach of adding points based on every t -tuple, $t < d$, as such a t -tuple of R_C might not even intersect C ! For use in constructing weak ε -nets whose size is low-dimensional, i.e., $o(1/\varepsilon^d)$, the set R_C must satisfy a condition of ‘closeness’ which is based on subsets of size lower than d . In fact, we can strengthen the statement to show that there exists an $R_C \subset R$ such that C intersects the convex hull of every $\lfloor d/2 \rfloor + 1$ -sized subset of R_C . Showing the existence of such a set R_C required generalizing a basic theorem of convex geometry, Carathéodory’s theorem. In fact that opens up a new set of generalizations of such theorems, to which we now turn.

★ ★ ★

⁹Helly’s theorem returns a point in the common intersection of a set of convex objects. A vertex of this common intersection will do, and it is the intersection of d planes, each plane defined by d points of P .

¹⁰In \mathbb{R}^2 , the descriptive complexity of approximate centerpoints can be shown to be at most 3 (rather than 4): take a vertical line l_1 through the median point (by x -coordinate) of P , and let l_2 be the line that equi-partitions the points on the left and right side of l_1 (such a line always exists by the Ham-sandwich theorem). Then observe that the point $l_1 \cap l_2$ is an approximate centerpoint, and is defined by 3 points of P .

A generalization of Carathéodory’s theorem. Our program leads us to study the behavior of *low-dimensional* simplices with respect to convex sets in \mathbb{R}^d . We will examine some classical theorems in discrete geometry – Radon’s theorem [Rad47], Carathéodory’s theorem [Mat02], colorful Carathéodory theorem [Bár82] – and prove extensions that demonstrate the phenomenon of low-dimensional intersections.

Carathéodory’s Theorem states¹¹ that if a convex set C intersects the convex hull of some point set P , then it also intersects some simplex spanned by points in P . Equivalently, either P can be separated from C by one hyperplane, or C intersects the convex hull of some $(d + 1)$ points of P . In general, this cannot be strengthened to guarantee lower-dimensional intersections, i.e., that the convex-hull of some $t < d + 1$ points of P intersect C . Though in \mathbb{R}^2 , it is easy to see that given a convex object C and a set P of points, either P can be separated from C using 2 lines, or there exist 2 points of P whose line-segment intersects C . In fact, we will show that a similar statement is true for \mathbb{R}^d – that by using more than one hyperplane to separate C from P , one can achieve low-dimensional intersections.

Theorem (Chapter 4). *Given a set P of n points in \mathbb{R}^d and a convex object C , either P can be separated from C by $O(d^4 \log d)$ hyperplanes (i.e., each $p \in P$ is separated from C by one of the hyperplanes), or C intersects the convex hull of some $\lfloor d/2 \rfloor + 1$ -sized subset of P .*

Note that apart from quantitative improvements on the number of separating hyperplanes, one cannot guarantee an even lower-dimensional intersection – to see this consider P to be the vertex set of a cyclic polytope and let C to be a slightly shrunk copy of the polytope.

The proof in Chapter 4 uses linear-programming duality, but the idea should be understood as an application of the beautiful multiplicative weights technique: given P and C , assume that the convex-hull of no $\lfloor d/2 \rfloor + 1$ -sized subset intersects C . Then it is not too hard to see that there must exist a halfspace h_1 containing a constant fraction of points of P (constant depends on d) and containing C in its complement. Now add to P an extra copy of each point of P not separated from C by h_1 . Repeat this process by picking another halfspace h_2 w.r.t. to the new pointset, and making an extra copy of each point of the new pointset not contained in h_2 and so on. One can visualize the making of the extra copies of p as doubling the *weight* of p (set to 1 initially). After some t steps, each point $p \in P$ must be contained in many of these t halfspaces: if p is contained in too few halfspaces, its weight would be very large (as it is doubled for each halfspace not containing p), contradicting the fact that the total weight of all the points in the end is not too much (as at each step we picked a halfspace containing a large weight of points in P). Now taking an ε -net of the halfspaces (i.e., the base set is $\{h_1, \dots, h_t\}$ and the sets are subsets of the base set containing a common point $q \in \mathbb{R}^d$) gives us a small-sized set of halfspaces separating all points of P from C .

A beautiful extension of Carathéodory’s theorem, the colorful Carathéodory Theorem, was discovered by Imre Bárány [Bár82]: given $d + 1$ sets of points P_1, \dots, P_{d+1} in \mathbb{R}^d and a convex set C such that $C \cap \text{conv}(P_i) \neq \emptyset$ for all $i = 1, \dots, d + 1$, there exists a set Q with $C \cap \text{conv}(Q) \neq \emptyset$ and where $|Q \cap P_i| = 1$ for all i . Equivalently, either some P_i can be separated from C with one hyperplane, or C intersects the convex hull of a rainbow set of $d + 1$ points¹².

¹¹Here we have stated the theorem in a slightly more general form; usually it is stated where C is just a point.

¹²This theorem is also commonly stated for the case where C is a point, but the above slight generalization follows immediately from Bárány’s proof technique [Bár82].

In \mathbb{R}^3 , an elementary argument shows that a similar low-dimensional generalization also holds here: given C and a set of red and blue points in \mathbb{R}^3 , either the red set or the blue set can be separated from C by a constant number of planes. Or there is a red-blue edge intersecting C . First, elementary considerations show that given a pointset and a convex set C , either there is a triangle Δ spanned by the points so that each edge of Δ intersects C or all the point can be separated from C using twelve hyperplanes. Now suppose that we have some red and blue points and a convex set C . Then applying the above result to each set of points, we conclude that either one of the sets can be separated from C with twelve planes or there is a red triangle and a blue triangle each of whose edges intersect C . For each vertex of these triangles consider the region on the boundary of C that it can see (imagine C to be opaque). Since each red (resp. blue) edge intersects C , no two of the red (resp. blue) regions intersect, i.e., no three of the six regions intersect at a common point. Since the regions are pseudodisks, their intersection graph is planar. As $K_{3,3}$ is not planar, there is a red region and a blue region which do not intersect. This implies that the red-blue edge defined by the points corresponding to these regions intersects C .

We will prove the lower-dimensional generalization of Colorful Carathéodory's Theorem in \mathbb{R}^d :

Theorem (Chapter 4). *For any d , there exists a constant N_d such that given $k = \lfloor d/2 \rfloor + 1$ sets of points P_1, \dots, P_k in \mathbb{R}^d and a convex object C , either one of the sets P_i can be separated from C by N_d hyperplanes, or there is a rainbow set of size k whose convex hull intersects C .*

We feel that this generalization is also true for several other basic theorems:

Conjecture 1 (Low-dimensional Kirchberger's Theorem). *Given a set P of n red points and n blue points in \mathbb{R}^d either there exist a constant number of hyperplanes \mathcal{H} such that every red-blue pair is separated by a plane in \mathcal{H} , or a $\lfloor d/2 \rfloor$ -dimensional red simplex intersects a $\lceil d/2 \rceil$ -dimensional blue simplex (or vice versa).*

Conjecture 2 (Low-dimensional Tverberg's Theorem). *Given a set P of n points in \mathbb{R}^d , there exists a Tverberg partition on a large-enough subset of P where two sets have size $\lfloor d/2 \rfloor + 1$.¹³*

* * *

The hitting-set problem. So far we have looked at the combinatorial problem of finding small-sized hitting sets for set-systems (X, \mathcal{F}) where each set has large size, at least $\varepsilon|X|$. It turns out that it is intimately related to the fundamental computational problem of computing small-sized hitting sets for arbitrary set-systems. For example, given n points P in the plane, and a set of m disks D , compute the smallest-sized subset $P' \subseteq P$ such that each disk in D contains at least one point of P' . This is the *geometric hitting-set problem* for disks in the plane (its dual problem is the *geometric set-cover problem* where one would like to pick the minimum number of disks $D' \subseteq D$ that cover all the points of P).

The geometric hitting-set problems are also normally NP-hard (even in the case where D are disks of the same radii [HM87]), so the best one can hope for is an efficient algorithm that returns

¹³It is not too hard to see that this implies a version of the extension of Carathéodory's theorem.

a $(1 + \epsilon)$ -approximation to the optimal solution, i.e., a polynomial-time approximation scheme (PTAS). Along this (still ongoing) quest, over the past three decades, there have been a number of algorithms proposed to compute good approximations to the optimal geometric hitting-set/set-cover for a variety of geometric objects (see [CKLT07, NV06, CMWZ04, CV07, AES12, CF13, CDD⁺09, HM84, NV06, EvL08] for a few examples). Much work remains to be done; for example, till recently a PTAS was not known for the case where the geometric objects are disks in the plane, the best approximation-factor being a 18-approximation algorithm [DFLON11]¹⁴.

In 1994, Bronnimann and Goodrich [BG95] showed the following surprising connection: let $\mathcal{R} = (P, \mathcal{D})$ be a set-system for which we want to compute a minimum hitting set. If one can compute an ϵ -net of size $f(1/\epsilon)/\epsilon$ for the weighted ϵ -net problem for \mathcal{R} in polynomial time then one can compute a hitting set of size at most $f(\text{OPT}) \cdot \text{OPT}$ for \mathcal{R} , where OPT is the size of the optimal hitting set, in polynomial time. A shorter, simpler proof was given by Even *et al.* [ERS05]. The connection is less surprising when one knows the proof: first assign weights to each point such that the total weight of the points in each set in \mathcal{D} is a large fraction of the total weight; then picking an ϵ -net (w.r.t. the weighted sets) hits every set! Assigning such weights is easy, as the optimal hitting-set gives such an assignment: $p \in P$ has weight 1 if p is in the optimal hitting set, 0 otherwise. The total weight is OPT , while each set has weight at least 1; so $\epsilon = 1/\text{OPT}$ and then ϵ -nets give a hitting set of size $f(\text{OPT}) \cdot \text{OPT}$. This of course requires knowing the optimal hitting set, but as the weights need not be integral, one can use the linear program for the problem to set the weights [ERS05]. Alternatively, multiplicative weights technique can be used to set the weights [BG95]. Note here the thematic use of ϵ -nets as *rounding* tools: the LP solution gives fractional weights, and ϵ -nets round these to integers by (essentially) picking a subset randomly with probability proportional to these weights.

Unfortunately, the Bronnimann-Goodrich technique cannot lead to a PTAS for geometric hitting set problems. This is a fundamental limitation of the technique: it *cannot* give better than constant-factor approximations. Apart from simple set-systems like intervals in \mathbb{R} , the sizes of ϵ -nets are at least c/ϵ for some constant $c > 1$, making it impossible to get below a c -approximation factor. We now show that a simple ‘local-search’ scheme can bypass the limitation of Bronnimann-Goodrich to get PTAS for several geometric set systems – arbitrary disks in the plane, halfspaces in \mathbb{R}^3 , pseudo-disks in the plane and several other basic systems:

Theorem (Chapter 6). *One can compute a $(1 + \epsilon)$ -approximation to the minimum-hitting set for set-system induced by disks in the plane, in time $O(n^{1/\epsilon^2})$.*

The algorithm starts with any hitting set $S \subseteq P$ (e.g., take all the points of P), and iterates local-improvement steps of the following kind: If any k points of S can be replaced by $k - 1$ points of P such that the resulting set is still a hitting set, then perform the swap to get a smaller hitting set. Halt if no such local improvement is possible. The main result is that this algorithm halts with a hitting-set of size at most $(1 + 1/\sqrt{k}) \cdot \text{OPT}$, where OPT is the size of the minimum hitting set¹⁵. In fact, the result is more general: it holds for any set of objects which are r -admissible. It also gives a similar PTAS for the hitting-set problem for halfspaces in \mathbb{R}^3 .

¹⁴That too for the special case where D is a set of unit disks.

¹⁵See the blog-post at the geometry blog GEOMBLOG for a nice exposition by Suresh Venkat of this technique: <http://geomblog.blogspot.com/2009/06/socg-2009-local-search-geometric.html>.

Such algorithms are generally classified as ‘local-search algorithms’, and have been used for decades as heuristics in fast practical methods. However widespread their use, it is often not easy to bound the accuracy of the resulting solutions. The fact that the final solution cannot be ‘locally’ improved does not normally imply closeness to the global optimum. Therefore theoretically there have been very few results demonstrating the mathematical soundness of many such search heuristics based on local search. Let us just mention the inspiration of our result, that of a surprising and beautiful use of local-search for the well-studied k -median problem to achieve a constant-factor approximation by Arya *et al.* [AGK⁺01], which was later simplified and extended to the k -means problem by Kanungo *et al.* [KMN⁺02]. A well-known example of a heuristic which works well in practice for clustering a set of points in \mathbb{R}^d is the so-called Lloyd’s method or k -means algorithm. The problem asks for computing k clusters minimizing the sum of square distances of the points to their cluster centers. A related problem asks for minimizing just the sum of distances to the k centers. Variant of this were studied in [5] for the case of *projective clustering*.

As a side result, Chapter 6 also shows that the local search technique can be used to prove the existence of small-sized ϵ -nets. Specifically, we show that for the case where we have points in the plane and ranges consist of unit squares in the plane, a simple local-search method gives the optimal bound of $O(1/\epsilon)$ for the size of the ϵ -net. It is quite easy to prove the same result using other techniques but it is interesting that the local search technique can be used to prove this. This kind of result is currently known only for half-spaces in \mathbb{R}^2 and is implied by the proof of the existence of $O(1/\epsilon)$ size ϵ -nets by Pach and Woeginger [PW90]. We conjecture that this holds for ϵ -nets w.r.t. halfspaces in \mathbb{R}^3 as well. Of course, given that *i*) halfspaces in \mathbb{R}^3 have linear-sized nets [MSW90, PR08] and *ii*) the hitting-set algorithm computes a $(1 + \epsilon)$ -approximation to the smallest-sized hitting set, we know that the local-search algorithm will produce a ϵ -net of size $O(1/\epsilon)$. The question is whether one can prove combinatorial bounds directly out of local-search, without first relying on existence results of small-sized nets.

* * *

The independent-set problem. We now illustrate the application of the local-search technique to the geometric independent-set problem, a basic optimization problem that is hopelessly difficult to approximate for general graphs. Recall the problem: given a set $\{O_1, \dots, O_n\}$ of n objects in the plane, the intersection graph has a vertex for each object, and an edge between two vertices if the two corresponding objects intersect. Computing a maximum independent set in this graph is equivalent to picking a maximum subset of disjoint objects.

Computing $\text{OPT}(S)$ is known to be *NP*-complete if S is a set of unit disks or a set of orthogonal segments in \mathbb{R}^2 [IA83]. For unit disks in \mathbb{R}^2 , a polynomial time $(1 + \epsilon)$ -approximation scheme was proposed in [HMR⁺98]. For arbitrary disks, independently Erlebach *et al.* [EJS01] and Chan [Cha03] presented a polynomial time $(1 + \epsilon)$ -approximation scheme. For the case of axis-parallel rectangles in the plane, Agarwal *et al.* [AvKS98] presented a $O(\log n)$ -approximation algorithm in time $O(n \log n)$, which was improved to a $O(\log \log n)$ -approximation algorithm [CC09]. Berman *et al.* [BDMR01] show that a $\log_k n$ -approximation can be computed in $O(n^k \alpha(S))$ time. For the case when the objects are arbitrary line segments in the plane, Agarwal-Mustafa [4] gave an algorithm

with approximation ratio $O(n^{1/2+o(1)})$. This was improved by Fox-Pach [FP11] to $O(n^\varepsilon)$, for any $\varepsilon > 0$, using small-sized separators in the intersection graphs of line segments.

We consider the case when the objects are axis-parallel rectangles in the plane, and show that when the independent-set has a large size, then a constant-factor approximation can be computed. The core of this algorithm illustrates again the use of local-search¹⁶ (on subsets of non-piercing rectangles): start with any independent set of rectangles, and check if one can improve the solution by replacing two rectangles from our current independent set with three rectangles. We will show that if one cannot make any such improvement, then the resulting independent set is a constant-factor approximation!

Theorem (Chapter 5). *Given a set of axis-parallel rectangles in the plane with maximum independent set of size βn for some $\beta < 1$, one can compute an independent set of size $\Omega(\beta^2 n)$ in polynomial time.*

The straightforward extension where k rectangles are replaced by $k+1$ gives a $(1 - O(1/\sqrt{k}))$ -factor approximation [CHP09]. After the appearance of the papers [4, 6, 29, CHP09], several works followed using this technique to derive PTAS for other geometric optimization problems. Gibson *et al.* [GKKV09] improved a previous 4-approximation algorithm to get a PTAS for a terrain guarding problem. Similar technique was also used to get a PTAS for dominating set in disk graphs [GP10].

* * *

Centerpoints. We have seen the useful role of centerpoints in various questions studied so far. In fact, centerpoints are just one measure of geometric data-depth. As we will see, there are several natural ways to measure data-depth which have been studied in the literature, related to each other in sometimes surprising ways. With each such measure there are two questions: *i*) proving the existence of a point which suitably captures, with some guaranteed bounds, that measure and *ii*) devising efficient algorithms to compute this point.

By now there are several proofs of the centerpoint theorem: using Brouwer’s fixed-point theorem, using Helly’s theorem, following from Tverberg’s theorem, by induction on the dimension d , using an elementary extremal argument and several others. Interestingly, each of these above techniques can be used to generalize this theorem in a different direction; the extremal-argument proof will be generalized later.

A particularly elegant one is as follows: replace ‘point’ by a ‘ball’ in the above definition; i.e., find a ball B such that any halfspace containing B contains at least $n/(d+1)$ points of P . This is trivial: any large-enough ball will do. Now it can be shown that B can be shrunk continuously to a point while maintaining this starting invariant! If B gets ‘stuck’ at some point in this shrinking process, then there must be $d+1$ halfspaces, each containing exactly $n/(d+1)$ points of P , such that *i*) their bounding planes are tangent to B and contain at least one point of P , and *ii*) the convex-hull of the $d+1$ intersections of these planes with B contains the center of B . Elementary geometric and counting considerations show that such a configuration is not possible. The essence of this

¹⁶In fact it is one of the first combinatorial algorithms that was able to provably show the effectiveness of local-search on such problems.

clever technique is: if one wants to prove that a point with a certain property exists, start with a ball satisfying the analogue of the property, and then show that the ball can be shrunk continuously to a point while satisfying the initial invariant. Tverberg-Vrecica proved this for Tverberg's theorem [TV93]. Simplification to the centerpoint theorem and a generalization to centerdisks was given in [10]. Helly's theorem has a similar proof. We conjecture that other theorems of similar type also have such proofs: e.g., the regression-depth theorem [ABET00], and the intersecting-rays theorem [FHP08].

Recall that the centerpoint theorem is equivalent to the fact that all convex-sets containing greater than $dn/(d+1)$ points of P can be hit by one point¹⁷; the generalization to hitting with k points is the weak ε -net problem for which only very partial results are known. Then, a first step towards the weak ε -net problem (and generalizing the centerpoint theorem) is to resolve the case when k is a small constant: what is the smallest parameter ε such that for any set $P \subset \mathbb{R}^2$, it is possible to find two points q_1 and q_2 such that any convex set containing $\varepsilon|P|$ points must contain either q_1 or q_2 ? In general, let ε_i^d be the minimum value of ε such that given any set P in \mathbb{R}^d , there exist i points hitting all convex objects containing greater than $\varepsilon|P|$ points of P . The centerpoint theorem can be restated as $\varepsilon_1^d = d/(d+1)$.

Aronov *et al.* [AAH⁺09] proved that given a set P of n points in the plane, all convex sets containing greater than $5n/8$ points of P can be hit by two points. They also construct inputs where regardless of how one picks the two points, there exists a convex set containing at least $5n/9$ points that is not hit. In our notation, $5/9 \leq \varepsilon_2^2 \leq 5/8$. They proved further results for larger number of points.

We improve the previous-best upper-bound of $5n/8$ to $4n/7$ for hitting with two points in the plane. For the same problem, we then improve the previous-best lower-bound of $5n/9$ to $4n/7$, thus completely resolving this problem. Our technique similarly improves the best-known bounds for hitting with 3, 4 and 5 points (improving the result of Alon-Rosenfeld [Mat02, p. 259]). The general upper-bound theorem we prove is:

Theorem (Chapter 7). *Given a set P of n points in \mathbb{R}^d , and integers $r, s \geq 0$,*

$$\varepsilon_{r+ds+1}^d \leq \frac{\varepsilon_r^d \cdot (1 + (d-1)\varepsilon_s^d)}{1 + \varepsilon_r^d \cdot (1 + (d-1)\varepsilon_s^d)}$$

where we define $\varepsilon_0^d = 1$.

Note that the centerpoint theorem is a consequence of this more general result; set $r = s = 0$! The optimal result for two points in the plane follows by $r = 1, s = 0, d = 2$. The proof uses an extremal argument that is a generalization of the argument used to prove Helly's theorem¹⁸ in \mathbb{R} : consider the interval I whose right endpoint is the leftmost of the right endpoints of all given intervals. This right endpoint must then hit all intervals, as any interval to the left violates the extremal choice of I while an interval to the right violates the pairwise-intersecting property. It turns out that this argument actually uses convexity, and so considering points over appropriately-defined extremal d -tuples works to prove Helly's theorem in \mathbb{R}^d . This is the basis for the proof of

¹⁷Clearly a centerpoint q must hit all convex-sets containing greater than $dn/(d+1)$ points, for otherwise the halfspace containing q and separating q from any such unhit convex-set contains less than $n/(d+1)$ points. Similarly, any point q hitting all such convex-sets must be a centerpoint.

¹⁸Any set of pairwise-intersecting intervals in \mathbb{R} can be hit by one point.

the above theorem. This extremal argument has been re-discovered and used in many contexts; e.g., the proof of colorful Helly's theorem by Lovász [Lov74] or Helly-type theorems for generalized linear-programming problems [Ame94].

Define the Tukey-depth [Tuk75] of a point q to be the minimum number of points contained in any half-space containing q . Then the centerpoint theorem can be re-phrased as: there always exists a point of Tukey-depth at least $n/(d+1)$. The point of highest Tukey-depth w.r.t. P is called the *Tukey median* of P , and its depth called the *Tukey depth* of P . In general, the set of points of Tukey depth at least βn form a convex region called the β -deep region of P . The β -deep region is non-empty for any $\beta \leq 1/(d+1)$. It is the intersection of all halfspaces containing more than $(1-\beta)n$ points of P . Each facet of this region is supported by a hyperplane that passes through d points of P .

★ ★ ★

Simplicial Depth. Observe that any centerpoint must be contained in many simplices spanned by $(d+1)$ -tuples of points in P . This leads to another related useful depth-measure: given a set P of n points in \mathbb{R}^d , the *simplicial depth* of a point q is the number of simplices spanned by points of P that contain q . The simplicial depth of P is the highest such depth of any point $q \in \mathbb{R}^d$.

The first question that arises is if there exists, for any P , a point q with high simplicial depth. A classic result of Bárány [Bár82] gave the first such combinatorial bound, showing in fact that there exists a point contained in at least $\frac{1}{d!(d+1)^{d+1}} \cdot n^{d+1}$ simplices spanned by P . Let c_d be a constant, depending on d , such that any pointset P has simplicial depth at least $c_d \cdot n^{d+1}$. The optimal dependency on the dimension, c_d , is a long-standing open problem. Bukh, Matoušek and Nivash [BMN10] showed an elegant construction of a point set P so that no point in \mathbb{R}^d is contained in, up to lower-order terms, more than $(n/(d+1))^{d+1}$ simplices defined by P . Furthermore, they conjecture that this is the right bound.


Conjecture 3 (Simplicial depth conjecture). *Given any set P of n points in \mathbb{R}^d , there always exists a point contained in at least $(n/(d+1))^{d+1}$ simplices spanned by P .*

For $d=2$, the above conjecture was solved in 1984 by Boros-Furedi [BF82] (see Bukh [Buk06] for an elegant proof). Yet another proof follows from the work of Fox *et al.* [FGL⁺11]. For $d=3$, the conjectured bound is $c_3 = 1/4^4 = 0.0039$.

To get some intuition for the statement (and warming up to an improvement), consider the following lemma (Chapter 8) of independent interest that gives a structural property of points as a function of their Tukey depth:

Lemma (Tukey Structural Lemma). *Given a set P of n points in \mathbb{R}^d , where $\text{depth}(P) = \tau n - 1$, there exists a point p with depth $\tau n - 1$, and a set \mathcal{H} of $d+1$ open halfspaces $\{h_1, \dots, h_{d+1}\}$, such that i) $|h_i \cap P| = \tau n$, ii) p lies on the boundary plane of each h_i , and iii) $h_1 \cup \dots \cup h_{d+1}$ covers $\mathbb{R}^d \setminus \{p\}$.*

To see the implications of this, consider the case when P is a set of n points in \mathbb{R}^2 . And the Tukey depth of P is $n/3$ (i.e., the minimum that is guaranteed by the centerpoint theorem). Then the lemma guarantees three halfspaces h_1, h_2, h_3 , each containing exactly $n/3$ points (see figure). Counting considerations imply that the three regions $h_1 \cap \bar{h}_2 \cap \bar{h}_3$, $\bar{h}_1 \cap h_2 \cap \bar{h}_3$ and $\bar{h}_1 \cap \bar{h}_2 \cap h_3$ each contain $n/3$ points. As any transversal contains the point q , there are $(n/3)^3$ simplices containing q .



Broadly speaking, the structure of this point-set is that of a triangle, where each vertex of the triangle has $n/3$ copies (as ‘enforced’ by the three planes). This structure works in any dimension: if the Tukey depth of a point-set is $n/(d + 1)$ then the points can be partitioned into $d + 1$ sets which behave like the vertices of a simplex. And, crucially, the simplicial depth conjecture can be seen to be true immediately! If it could be shown that the simplicial-depth is an increasing function of Tukey-depth, this would then imply a positive solution to the simplicial depth conjecture in \mathbb{R}^d .

As the Tukey-depth of a point increases, the structure implied by the lemma becomes messier: points start ‘leaking’ from the $d+1$ regions into other regions, and transversals could stop containing q . However, it can be shown that this leaking process can be quantitatively controlled as a function of Tukey-depth, and each leaked point creates new simplices containing q .

Theorem (Chapter 8). *Any set P of n points in \mathbb{R}^3 has simplicial depth at least $0.0023 \cdot n^4$.*

In fact, experimental evidence reveals that the point computed in the above theorem satisfies the bounds of the conjecture:

Conjecture 4. *The point computed in the above theorem is contained in at least $(n/4^4)$ simplices spanned by P .*

Independently, using algebraic topology machinery, Gromov [Gro10] improved the bound to the value $c_3 \geq 0.0026$. This bound for \mathbb{R}^3 has since been improved even further by Matousek and Wagner [MW10] to 0.00263, and then by Kral *et al.* [KMS12] to 0.00031. All these methods use algebraic topology related machinery. Our proof, on the other hand, is elementary (though not the best) and we conjecture that it leads to the resolution of the conjecture. Gromov’s result in fact proves the bound for general d , showing that $c_d \geq 2d/((d + 1)!^2(d + 1))$. His proof has since been simplified by Karasev [Kar12].

* * *

Ray-Shooting Depth. Recently, an elegant new result has been discovered in \mathbb{R}^2 that easily implies both the centerpoint theorem and the simplicial depth theorem in the plane. Given P , let E be the set of all $\binom{n}{2}$ edges spanned by points of P . Then a point $q \in \mathbb{R}^2$ has *Ray-Shooting depth* at

least r if any half-infinite ray from q in any direction $u \in \mathbb{S}^1$ intersects at least r edges in E . The Ray-Shooting depth (henceforth called RS-depth) of P is the maximum RS-depth of any point in \mathbb{R}^2 . Fox-Gromov-Lafforgue-Naor-Pach [FGL⁺11] proved the following statement:

Theorem (Fox *et al.* [FGL⁺11]). *Given a set P of n points in \mathbb{R}^2 , there exists a point q with RS-depth at least $n^2/9$.*

Note that this statement immediately implies both the centerpoint theorem and the simplicial-depth theorem in the plane¹⁹. The proof in [FGL⁺11] is topological, and does not give a method to compute such a point. Efficient algorithms were obtained in Mustafa-Ray-Shabbir [31] by ‘de-topologising’ their proof to a combinatorial one:

Theorem (Mustafa *et al.* [31]). *Given a set P of n points in \mathbb{R}^2 , a point of RS-depth $\Omega(n^2/9)$ can be computed in time $O(n^2 \log^2 n)$.*

It turns out that RS-depth immediately implies improved results for another depth measure. The Oja depth [Oja83] of a point $x \in \mathbb{R}^d$ w.r.t. P is defined to be the sum of the volumes of all $(d+1)$ -simplices spanned by x and d points of P , normalized with respect to the volume of the convex hull of P ²⁰. The Oja depth of P , denoted $\text{Oja-depth}(P)$, is the minimum Oja depth over all $x \in \mathbb{R}^d$. Construct P by placing $n/(d+1)$ points at each of the $(d+1)$ vertices of a unit-volume simplex in \mathbb{R}^d . It is easy to see that any point will have Oja depth at least $(n/(d+1))^d$. The conjecture in [CDI⁺13] is that the lower bound given above is tight.

Conjecture 5 (Oja-depth conjecture). *For all sets $P \subset \mathbb{R}^d$ of n points, $\text{Oja-depth}(P) \leq (\frac{n}{d+1})^d$.*

For general d , the previous best upper bound [CDI⁺13] was that the Oja depth of any set of n points is at most $\binom{n}{d}/(d+1)$. For $d=2$, this is $n^2/6$. This can immediately be improved by realizing a connection to RS-depth: let q be the point realizing RS-depth of $n^2/9$ for P . Then the Oja-depth of q is at most $n^2/7.2$: the number of triangles (spanned by pairs of points in P and the point q) containing any point $p \in \mathbb{R}^2$ is at most the number of edges spanned by P intersecting the ray \vec{qp} , which is at most $n^2/4 - n^2/9 = n^2/7.2$. Integrating over all p gives the required bound. A more complete calculation in \mathbb{R}^d gives an improvement to the previous-best bound by orders of magnitude:

Theorem (Chapter 9). *Every set P of n points in \mathbb{R}^d , $d \geq 3$, has Oja depth at most*

$$\frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2 (d+1)!} \binom{n}{d} + O(n^{d-1}).$$

We also resolve this conjecture in two dimensions:

Theorem (Chapter 9). *Given a set P of n points in \mathbb{R}^2 , $\text{Oja-depth}(P) \leq n^2/9$.*

¹⁹Let q be a point with RS-depth $n^2/9$. Consider any line l through q . Then l must intersect at least $2n^2/9$ edges, and therefore both halfspaces defined by l must contain at least $(n - n\sqrt{1-8/9})/2 = n/3$ points. For simplicial depth, consider, for each point $p \in P$, the ray from q in the direction \vec{pq} . Then for every edge $\{p_i, p_j\}$ that intersects this ray, the triangle defined by $\{p, p_i, p_j\}$ must contain q . Summing up these triangles over all points, each triangle can be counted three times, and so q lies in at least $n^3/27$ distinct triangles.

²⁰So assume w.l.o.g. that $\text{vol}(\text{conv}(P)) = 1$.

The Oja depth conjecture states the existence of a low-depth point, but given P , computing the lowest-depth point is also an interesting problem. In \mathbb{R}^2 , Rousseeuw and Ruts [RR96] presented a $O(n^5 \log n)$ time algorithm for computing the lowest-depth point, which was then improved to the current-best algorithm with running time $O(n \log^3 n)$ [ALST03]. For general d , various heuristics for computing points with low Oja depth were given by Ronkainen, Oja and Orponen [ROO03]. Finally, Krishnan-Mustafa-Venkatesh [16] presented an approximate algorithm utilizing fast rendering systems on current graphics hardware for computing Oja-depth contours in the plane.

★ ★ ★

A Unifying View. I think that one good way to answer questions related to geometric data-depth is by the following analogy. Construct a point set of size n by fixing a simplex in \mathbb{R}^d (does not have to be regular) and placing $n/(d+1)$ points at each of its $(d+1)$ vertices. Call such a point set a *simplex-like* point set. For questions related to data-depth, a simplex-like set seems to represent the worst case: if some data-depth property is true for this point set, then it is true for any point set.

I mention two facts in the support of this intuition. First consider all depth measures we have examined thus far for the simplex-like point set (with the underlying simplex S). Take, say, the centroid c of this simplex. Then the centerpoint theorem follows for this point set because any halfspace containing c must contain at least one vertex of S , and so contains $n/(d+1)$ points. Similarly, taking one point from each of the vertices defines a simplex containing c , and so c is contained in $(n/(d+1))^{d+1}$ simplices. Finally, any ray from c must intersect at least one facet of S , and so has RS-depth $(n/(d+1))^d$. The intuition one gets from simplex-like point sets is in accordance with every information we know about these problems.

Second, as we have already outlined earlier, when the Tukey-depth of P is the lowest possible, i.e., $n/(d+1)$, P behaves like a simplex-like pointset: the $d+1$ halfspaces $\mathcal{H} = \{h_1, \dots, h_{d+1}\}$ specified by the Tukey Structural Lemma are such that the $d+1$ regions $A_i = h_i \cap (\bigcap_{j \neq i} \bar{h}_j)$ each contain exactly $n/(d+1)$ points, and all the other $2^{d+1} - 2$ regions are empty. And it is not too hard to prove that in such a configuration in \mathbb{R}^d , the point p has Tukey-depth $n/(d+1)$, it has simplicial-depth $(n/(d+1))^{d+1}$, and has RS-depth $(n/(d+1))^d$. This, together with the first point, leads us to suspect that the bounds derived from the simplex-like point set might indeed be always realizable for any point set.

This viewpoint opens up the possibility of the existence of a much broader structure for point sets in \mathbb{R}^d – integrating centerpoints, simplicial depth and RS-depth into a uniform hierarchy of depth-measures. For $d = 1$, we have two measures: Tukey-depth and Simplicial-depth, defined by a 1-dimensional half-flat (i.e., a halfspace in 1D) and a 0-dimensional half-flat respectively. For $d = 2$, we now know of three measures: Tukey-depth, RS-depth and Simplicial-depth, defined by 2- and 1- and 0-dimensional half-flats respectively. Now consider the case $d = 3$. And let $c \in \mathbb{R}^3$. Then by considering the 3-dimensional space defined by a halfspace with c on its (2-dimensional) boundary, we get the notion of Tukey depth. By considering the 1-dimensional space defined by a half-line with c on its (0-dimensional) boundary, we get the notion of RS-depth. But this begs the question: what about 2-dimensional space with c on its (1-dimensional) boundary? The natural answer is to consider the 2-dimensional space defined by a half-plane h with c on its (1-dimensional) boundary.

And then count the number of edges spanned by P that intersect h .

Formally, a point $q \in \mathbb{R}^3$ has *Line-depth* r if any halfplane through q intersects at least r edges spanned by P . The Line-depth of a point set P is the highest Line-depth of any point. The simplex-like pointset has line-depth at least $(n/4)^2$, as any halfplane through c must intersect at least one edge of S . We conjecture that any set P of n points in \mathbb{R}^3 has Line-depth at least $(n/4)^2$; however we can show:

Theorem (Mustafa et al. [32]). *Given any set P of n points in \mathbb{R}^3 , there exists a point q such that any halfplane through q intersects at least $n^2/24.5$ edges spanned by P .*

Like RS-depth and Simplicial-depth in \mathbb{R}^3 , it seems hard to prove this exact bound using current techniques. Intuitively, it is clear that as the dimension of the flat decreases, the degrees of freedom increase and the problem becomes more complicated. On one end, optimal results for the 2-dimensional case (Tukey-depth) are known in any dimension. And on the other end, very partial results are known for the 0-dimensional case. It is our hope that the middle 1-dimensional case (line-depth) will be more within current reach than the 0-dimensional case. Based on our results on existence of points with high line-depth in \mathbb{R}^3 , we in fact can conjecture a ‘spectrum’ of structures:

Conjecture 6 (Centerflat Conjectures). *Given a set P of n points in \mathbb{R}^d , and an integer $0 \leq k \leq d$, there exists a point $q \in \mathbb{R}^d$ such that any $(d - k)$ -half flat through q intersects at least $(n/(d + 1))^{k+1}$ k -simplices spanned by P .*

The case $k = 0$ is the centerpoint theorem in any \mathbb{R}^d . The case $k = d$ is the simplicial-depth conjecture in any \mathbb{R}^d . The case $d = 2, k = 1$ is the RS-depth result of [FGL⁺11]. The case $d = 3, k = 1$ is on Line-depth.

Asymptotically tight (as a function of n) bounds for the centerflat conjecture follow naturally, though the precise dependence on d remains open (probably difficult, as even a special case of these conjectures, simplicial depth, is still open, even in \mathbb{R}^3).

Theorem (Mustafa et al. [32]). *Given a set P of n points in \mathbb{R}^d , and an integer $0 \leq k \leq d - 1$, there exists a point $q \in \mathbb{R}^d$ such that any $(d - k)$ -half flat through q intersects at least*

$$\max \left\{ \binom{n/(d+1)}{k+1}, \frac{2d}{(d+1)(d+1)! \binom{n}{d-k}} \cdot \binom{n}{d+1} \right\}$$

k -simplices spanned by P .

The proof follows from the use of Tverberg’s Theorem to partition P into $t = n/(d+1)$ sets P_1, \dots, P_t such that there exists a point q with $q \in \text{conv}(P_i)$ for all i . Consider any $(d - k)$ -dimensional half-flat \mathcal{F} through q , where $\partial\mathcal{F}$ is a $(d - k - 1)$ -dimensional flat containing q . Project \mathcal{F} onto a $(k + 1)$ -dimensional subspace \mathcal{H} orthogonal to $\partial\mathcal{F}$ such that the projection of \mathcal{F} is a ray r in \mathcal{H} , and $\partial\mathcal{F}$ and q are projected to the point q' . And let P'_1, \dots, P'_t be the projected sets whose convex-hulls now contains the point q' . Then note that the k -dimensional simplex spanned by $(k + 1)$ points $Q' \subset P'$ intersects the ray r if and only if the k -dimensional simplex defined by the corresponding set Q in \mathbb{R}^d intersects the flat \mathcal{F} . Now apply the single-point version ²¹ of Colorful Carathéodory’s

²¹Given any point $s \in \mathbb{R}^d$ and d sets P_1, \dots, P_d in \mathbb{R}^d such that each $\text{conv}(P_i)$ contains the origin, there exists a d -simplex spanned by s and one point from each P_i which also contains the origin.

Theorem [Bár82] to every $(k + 1)$ -tuple of sets, say P'_1, \dots, P'_{k+1} , together with the point s at infinity in the direction antipodal to the direction of r to get a ‘colorful’ simplex, defined by s and one point from each P'_i , and containing q' . Then the ray r must intersect the k -simplex defined by the $(k + 1)$ points of P' , and so the corresponding points of P in \mathbb{R}^d span a $(k + 1)$ -simplex intersecting \mathcal{F} . In total, we get $\binom{n/(d+1)}{k+1}$ k -simplices intersecting \mathcal{F} . Another way is to use the result of Gromov [Gro10], that given any set P of n points in \mathbb{R}^d , there exists a point q lying in $2d/((d + 1)(d + 1)!) \cdot \binom{n}{d+1}$ d -simplices. Now take any $(d - k)$ -half flat through q . It must intersect at least one k -simplex of each d -simplex containing it, and where each k -simplex is counted at most $\binom{n}{d-k}$ times. And we get

$$\frac{2d}{(d + 1)(d + 1)! \binom{n}{d-k}} \cdot \binom{n}{d + 1}$$

distinct k -simplices intersecting any $(d - k)$ -half flat through q .

Coming back to the simplex-like point set, one can further observe something stronger: take any line l through c and move l in any way to ‘infinity’ (i.e., outside the convex-hull of P). Then it still has to intersect at least one edge of the tetrahedron; i.e., the property is in fact topological in nature. This is already true for centerpoints²². We conjecture the same is true for \mathbb{R}^d :

Conjecture 7 (Topological centerflat Conjectures). *Given a set P of n points in \mathbb{R}^d , and an integer $0 \leq k \leq d$, there exists a point $q \in \mathbb{R}^d$ such that any $(d - k - 1)$ -flat through q must cross at least $(n/(d + 1))^{k+1}$ k -simplices spanned by P to move to infinity (i.e., so that the flat does not intersect the convex-hull of P).²³*

We end this overview by noting that these conjectures are true in \mathbb{R}^2 :

Theorem (Mustafa *et al.* [32]). *For any set P of n points in \mathbb{R}^2 , the k -centerpoint conjectures are true.²⁴*

²²If a point q has Tukey-depth r , then any plane through q has to cross at least r points to reach the point at infinity, regardless of whether this movement is any arbitrary continuous movement or only affine. See [ABET00].

²³The case $k = d$ gives a ‘-1’-flat moving to infinity, which we will treat as a stationary point.

²⁴ It has been communicated to us by a very insightful reader that this was implicit in Gromov’s paper [Gro10]. However the algebraic topology techniques there are highly non-trivial while the short proof in [32] is elementary.

This page was intentionally left blank

Part I

Around ε -nets

2. Geometric Separators

If one is trying to maximize the size of some structure under certain constraints, and if the constraints seem to force the extremal examples to be spread about in a uniform sort of way, then choosing an example randomly is likely to give a good answer.

Timothy Gowers

In this chapter we present the following two theorems. The first theorem improves the result in [AW13, AW14] as stated earlier, and shows the existence of near-optimal separators for weighted disjoint line segments:

Theorem 2.1. *Given a set S of n weighted disjoint line segments in the plane (with total weight W , and no segment having weight more than $W/3$), and a parameter $\delta > 0$, there exists a piecewise linear simple closed curve \mathcal{C} with $O(1/\delta)$ vertices such that *i*) the total weight of segments intersecting \mathcal{C} is at most δW , and *ii*) the total weight of the segments completely inside or outside \mathcal{C} is at most $2W/3$. Furthermore this is near-optimal, in the sense that any \mathcal{C} satisfying these two properties must have $\Omega(1/\delta)$ vertices.*

For the case where the segments are unweighted but not necessarily disjoint, one can show the existence of a more general separator (slightly generalizing the separator result of Fox-Pach [FP09] for the case of line segments):

Theorem 2.2. *Given a set S of n line segments in the plane with m intersections, and a parameter r , there exists a piecewise linear simple closed curve \mathcal{C} in the plane such that the number of segments completely inside (or outside \mathcal{C}) is at most $2n/3$ (call any such curve balanced), and*

- *the number of vertices of \mathcal{C} are $O\left(\sqrt{r + \frac{mr^2}{n^2}}\right)$, and*
- *the number of line-segments in S intersecting \mathcal{C} are $O\left(\sqrt{\frac{n^2}{r} + m}\right)$.*

Throughout the rest of this chapter, we will assume S to be a set of n weighted line segments in the plane. Let w_i denote the weight of segment $s_i \in S$, and W be the total weight. We assume S to be in general position, so no three line segments intersect at the same point.

A Partitioning Statement

The proof of both these statements follow from a suitable subdivision of the plane, and the application of a variant of the planar graph separator theorem. Our proof can also be seen as a generalization of the separator theorem of Fox-Pach [FP11] where, given a set of curves with m intersections, they show the existence of a separator that intersects $O(\sqrt{m})$ curves: this is obtained by applying the planar graph separator theorem on the arrangement induced by these curves (where each intersection is taken as a vertex). We also apply the planar graph separator theorem, but instead on a coarser subdivision of the plane. This subdivision is similar to a structure for the case of lines in the plane, called *cuttings* [Mat92]. We will assume the segments in S are in general position; in particular no three intersect at the same point.

The key statement from which both the separator theorems follow is:

Lemma 2.3. *Given a set S of n line segments in the plane with m intersections, and a parameter r , there exists a partition of \mathbb{R}^2 into $O(r + \frac{mr^2}{n^2})$ triangles such that the interior of any region in this partition intersects $O(n/r)$ segments.*

Proof. We first briefly review the basic partitioning method of using *trapezoidal decompositions*. Given a set $R \subseteq S$ of line segments, one can partition the space (say inside a large-enough rectangle containing all the segments of S) as follows. For each endpoint of a segment in R or an intersection-point between segments in R , shoot a vertical ray upwards (and downwards) till it hits another segment (or the bounding rectangle). The union of all these segments together with R partitions the bounding rectangle into a set of trapezoids (and triangles) called its *trapezoidal decomposition*¹. Denote by $\Xi(R)$ this set of trapezoids in the trapezoidal decomposition of R . The size, $|\Xi(R)|$, of the trapezoidal decomposition of R is the number of trapezoids in $\Xi(R)$; it is, within a constant-factor, equal to the total number of end- and intersection- points in R . A trapezoid present in the trapezoidal decomposition of any subset R of S is called a *canonical trapezoid*. For a canonical trapezoid Δ , let $|\Delta|$ denote the set of segments of S intersected by Δ . A crucial fact is that each trapezoid Δ is determined by a constant (3 or 4) number of segments in S . A trapezoid Δ is present in the trapezoidal decomposition of R if and only if its determining segments are present in R , and R does not contain any of the segments of S that intersect Δ . For the rest of the proof, we only work with canonical trapezoids determined by 4 segments. The case for canonical trapezoids determined by 3 segments is similar.

The proof is by an application of the ‘sampling refinement’ technique for constructing ε -nets. First note that a slightly weaker bound (within logarithmic factors) follows immediately from ε -nets. Given S , consider the set-system (S, \mathcal{F}) induced by intersection with segments in the plane, i.e.,

$$F \in \mathcal{F} \text{ iff there exists a line segment } l \text{ s.t. } F = \{s \in S \mid s \cap l \neq \emptyset\}$$

Pick a random set R by uniformly adding each segment of S with probability $p = (Cr \log r)/n$, where C is a large constant. Then R is a $(1/r)$ -net for (S, \mathcal{F}) with probability at least $9/10$. The expected size of R is np , and the expected number of intersections of segments in R is mp^2 . By Markov’s inequality, with probability at least $9/10$, the size of R is at most $10np$, and the number

¹We refer the reader to [Mat02] for a nice exposition on trapezoidal decompositions.

of intersections in R is at most $10mp^2$. Therefore with probability at least $8/10$, R is a $(1/r)$ -net and the size of the trapezoidal decomposition of R is $O(r \log r + (mr^2 \log^2 r)/n^2)$. Note that any open line-segment l in this trapezoidal decomposition must intersect at most n/r segments of S , as otherwise the set of segments intersecting l would not be hit by a segment from R , contradicting the fact that R is a $(1/r)$ -net. By dividing each trapezoid of $\Xi(R)$ into two triangles, we get a partition into $O(r \log r + (mr^2 \log^2 r)/n^2)$ triangles where the interior of each triangle intersects at most $O(n/r)$ segments of S .

The way to improve the above construction is by the so-called ‘‘sampling refinement’’ technique in the study of ε -nets. Set $p = Cr/n$ (for a small-enough constant C to be set later), and pick each segment in S with probability p to get a random sample R . Construct the trapezoidal decomposition $\Xi(R)$ of R . If all trapezoids $\Delta \in \Xi(R)$ intersect at most n/r segments in S , we are done. Otherwise we will further partition each violating Δ , based on two ideas. First, the expected number of trapezoids in $\Xi(R)$ intersecting more than n/r segments are few. In particular, we will show (Lemma 2.5) that the expected number of trapezoids intersecting at least tn/r segments in S is exponentially decreasing as a function of t . Second, consider a Δ intersecting a set, say S_Δ , of $n_\Delta = tn/r$ segments of S . Use the weaker bound on S_Δ with parameter t to get a partition inside Δ of $O(t \log t + (m_\Delta t^2 \log^2 t)/n_\Delta^2) = O(t^2 \log^2 t)$ triangles. By definition, each such triangle intersects at most $n_\Delta/t = n/r$ segments of S_Δ (and hence of S). Thus refining each Δ gives the required partition on S with parameter r . It remains to bound the overall expected size of this partition.

Lemma 2.4. *Given a set S of n line segments in the plane with m intersections, the number of canonical trapezoids defined by S that intersect at most k segments of S is $O(nk^3 + mk^2)$.*

Proof. Let $\Xi_{\leq k}$ be the set of canonical trapezoids defined by S that intersect at most k segments of S . The proof is standard via the Clarkson-Shor technique. Construct a sample T by adding each segment of S with probability p_0 ; the expected total number of picked segments is np_0 and the expected number of intersections between the segments of T is mp_0^2 . The trick is to count the expected size of $\Xi(T)$ in two ways. On one hand, it is at most $O(np_0 + mp_0^2)$ (i.e., the expected number of vertices present in $\Xi(T)$). On the other hand, as the probability of a canonical trapezoid Δ being in $\Xi(T)$ is $p_0^4(1 - p_0)^{|\Delta \cap S|}$, it is at least

$$\sum_{\Delta \in \Xi_{\leq k}} p_0^4(1 - p_0)^{|\Delta \cap S|} \geq \sum_{\Delta \in \Xi_{\leq k}} p_0^4(1 - p_0)^k$$

where the sum is over all canonical trapezoids Δ which intersect at most k segments of S . Therefore,

$$\begin{aligned} \sum_{\Delta} p_0^4(1 - p_0)^k &= |\Xi_{\leq k}| \cdot p_0^4(1 - p_0)^k \leq E[|\Xi(T)|] = np_0 + mp_0^2 \\ |\Xi_{\leq k}| &\leq \frac{np_0 + mp_0^2}{p_0^4(1 - p_0)^k} = O(nk^3 + mk^2) \end{aligned}$$

for $p_0 = 1/2k$. □

Lemma 2.5. *Expected number of trapezoids in $\Xi(R)$ intersecting at least tn/r lines of S is*

$$O\left(\left(t^3 r + \frac{mr^2 t^2}{n^2}\right)e^{-t}\right)$$

Proof. By definition:

$$E[|\Delta \in \Xi(R) \text{ s.t. } |\Delta \cap S| = tn/r|] = |\Delta \text{ s.t. } |\Delta \cap S| = tn/r| \cdot p^4(1-p)^{tn/r}$$

Using Lemma 2.4,

$$\begin{aligned} E[|\Delta \in \Xi(R) \text{ s.t. } |\Delta \cap S| = tn/r|] &\leq O(n(tn/r)^3 + m(tn/r)^2) p^4(1-p)^{tn/r} \\ &= O\left(t^3 r + \frac{mr^2 t^2}{n^2} e^{-t}\right) \end{aligned}$$

The bound follows by summing up over all trapezoids intersecting at least tn/r segments in S . \square

Now we can complete the proof of the theorem. Let $n_\Delta = t_\Delta n/r$ be the number of line segments in S intersected by each trapezoid $\Delta \in \Xi(R)$ (and m_Δ the number of their intersections). Using the weaker bound, refine trapezoid Δ by adding a $(1/t_\Delta)$ -net R_Δ for all the $t_\Delta n/r$ line segments of S intersected by Δ . The resulting expected total size of the trapezoidal partition is:

$$\begin{aligned} &= |R| + \sum_{\Delta} Pr[\Delta \in \Xi(R)] \cdot \text{Size of trapezoidal decomposition of } (1/t_\Delta)\text{-net within } \Delta \\ &= |R| + \sum_{\Delta} Pr[\Delta \in \Xi(R)] \cdot O\left(t_\Delta \log t_\Delta + \frac{m_\Delta t_\Delta^2 \log^2 t_\Delta}{n_\Delta^2}\right) \text{ (using the weaker bound)} \\ &\leq |R| + \sum_{\Delta} Pr[\Delta \in \Xi(R)] \cdot O(t_\Delta^2 \log^2 t_\Delta) \text{ (as } m_\Delta \leq n_\Delta^2) \\ &= |R| + \sum_j \sum_{\substack{\Delta \text{ s.t.} \\ 2^j \leq t_\Delta \leq 2^{j+1}}} Pr[\Delta \in \Xi(R)] \cdot O(t_\Delta^2 \log^2 t_\Delta) \\ &\leq |R| + \sum_j E[\text{Number of trapezoids } \Delta \text{ in } \Xi(R) \text{ with } 2^j \leq t_\Delta] \cdot O\left(2^{2(j+1)} \log^2 2^{j+1}\right) \\ &\leq |R| + \sum_j O\left(2^{3j} r + \frac{mr^2 2^{2j}}{n^2}\right) e^{-2^j} \cdot O\left(2^{2(j+1)} \log^2 2^{j+1}\right) \text{ (Lemma 2.5)} \\ &= |R| + r \sum_j O\left(2^{3j} e^{-2^j}\right) \cdot O\left(2^{2(j+1)} \log^2 2^{j+1}\right) + \frac{mr^2}{n^2} \sum_j O\left(2^{2j} e^{-2^j}\right) \cdot O\left(2^{2(j+1)} \log^2 2^{j+1}\right) \\ &= np + mp^2 + O(r) + O\left(\frac{mr^2}{n^2}\right) = O\left(r + \frac{mr^2}{n^2}\right) \text{ (the summands form a geometric series)} \end{aligned}$$

as required. This finishes the proof of Lemma 2.3.

Proof of Theorem 2.1

Given a set S of n weighted disjoint line segments of total weight W , in the standard way, by scaling, one can assume the weights are integral. Then make w_i copies of the segment $s_i \in S$ to get a set S' of W segments, each of weight 1. Apply Lemma 2.3 to S' ; if the segments in S' are

disjoint, the proof of Lemma 2.3 shows the existence of a subset R' of S' of total size $O(r)$ such that the trapezoidal decomposition of R' gives a partition \mathcal{T} where the interior of any region in \mathcal{T} intersects $O(W/r)$ segments of S' . Remove copies of the same segment in R' to get a subset R of S of total size $O(r)$. Now replace each $s \in R$ with a small triangle which contains s in its interior and modify \mathcal{T} accordingly to use edges of this new region instead of s . As each segment $s \in S$ of weight $\Omega(W/r)$ must be present in R (otherwise it would intersect the interior of some region of \mathcal{T} contradicting the partitioning property), this ensures that the new triangulation \mathcal{T}' has the properties that *i*) every $s \in S$ with weight $\Omega(W/r)$ lies in the interior of a single face of \mathcal{T}' , *ii*) the number of vertices of \mathcal{T}' is $O(r)$, and *iii*) each edge in \mathcal{T}' intersects segments in S of total weight $O(W/r)$ (for small-enough replacing triangles around each $s \in R$).

\mathcal{T}' can be seen as an embedding of an underlying planar graph G . Give weights to each face of \mathcal{T}' : if a segment $s \in S$ intersects t faces of \mathcal{T}' , add weight w_s/t to the weight of each of these t faces. A variant of the planar graph separator theorem [Mil86] now implies the existence of a simple cycle \mathcal{C} in \mathcal{T}' of $O(\sqrt{r})$ vertices such that faces completely inside (and outside) have total weight at most $2W/3$, and hence so do the segments of S inside (and outside) \mathcal{C} . The weight of the segments of S intersected by \mathcal{C} is at most $O(\sqrt{r}) \cdot O(W/r) = O(W/\sqrt{r})$. Setting $r = 1/\delta^2$ concludes the upper-bound.

The optimality of this statement can be seen by the following construction where S consists of a set of disjoint line segments of weight 1. Take a regular polygon \mathcal{P} with c/δ vertices, and place $\delta n/c$ copies of \mathcal{P} concentrically, each shrunk slightly more than the previous one so that there are no intersections between any two copies. Note that one can choose the scaling factor small-enough such that any closed curve separating two different copies of \mathcal{P} must also have at least c/δ vertices. Finally replace each polygon with c/δ line segments corresponding to its sides (slightly perturbed so that they are disjoint). Take any balanced closed curve \mathcal{C}' in the plane. If it contains at least one copy of \mathcal{P} completely inside, and one copy completely outside, then by construction it has at least c/δ vertices. Otherwise, say there is no copy of \mathcal{P} completely inside \mathcal{C}' . As \mathcal{C}' is balanced, it contains at least $n/3$ segments inside or intersecting its boundary; these segments belong to at least $(n/3)/(c/\delta) = \delta n/(3c)$ different copies of \mathcal{P} , and each of these copies must intersect \mathcal{C}' in at least one segment.

Proof of Theorem 2.2

Given the set S of n line segments with m intersections, apply Lemma 2.3 to get a partition \mathcal{T} of \mathbb{R}^2 into $O(r + mr^2/n^2)$ regions. \mathcal{T} can be seen as an embedding of an underlying planar graph G . Give weights to each face of \mathcal{T} : if a segment $s \in S$ intersects t faces of \mathcal{T} , add weight $1/t$ to the weight of each of these t faces. Now from [Mil86] we get a simple cycle \mathcal{C} in \mathcal{T} of $O(\sqrt{r + mr^2/n^2})$ vertices such that faces completely inside (and outside) have total weight at most $2n/3$, and hence so do the segments of S inside (and outside) \mathcal{C} . The number of segments of S intersected by \mathcal{C} is at most $O(\sqrt{r + mr^2/n^2}) \cdot O(n/r) = O(\sqrt{m + n^2/r})$.

This page was intentionally left blank

3. The General Weak ε -net problem

In dealing with mathematical problems, specialization plays, as I believe, a still more important part than generalization. Perhaps in most cases where we seek in vain the answer to a question, the cause of the failure lies in the fact that problems simpler and easier than the one in hand have been either not at all or incompletely solved. All depends, then, on finding out these easier problems, and on solving them by means of devices as perfect as possible and of concepts capable of generalization. This rule is one of the most important levers for overcoming mathematical difficulties and it seems to me that it is used almost always, though perhaps unconsciously.

David Hilbert

In this chapter we consider the general weak ε -net problem, and show that a random sample of size $O(1/\varepsilon \log 1/\varepsilon)$ can be used to construct a weak ε -net for P .

We first present an elementary proof for the two-dimensional case in Section 3.1. While this gives the intuition for the problem, the proof uses planarity strongly, and so the extension to higher dimensions uses a different approach based on the Hadwiger-Debrunner theorem. The general approach can be improved for \mathbb{R}^3 with additional ideas, which are presented in Section 3.2. The general construction for arbitrary dimensions is then presented in Section 3.3.

3.1 Two Dimensions

Consider the range space $\mathcal{R}_k = (P, R)$, where P is a set of n points in the plane, and $R = \{P \cap \bigcap_{i=1}^k h_i, h_i \text{ is any halfspace}\}$ are the subsets induced by the intersection of any k half-spaces in the plane. This range space has constant VC-dimension (depending on k), and from Welzl and Haussler [HW87], it follows that a random sample of size $O(1/\varepsilon \log 1/\varepsilon)$ is an ε -net for \mathcal{R}_k with some constant probability. Let Q be such a ε -net. We have the following structural claim which establishes a relation between strong ε -nets and weak ε -nets.

Lemma 3.1. *Let P be a set of n points in the plane, and let Q be an ε -net for the range space \mathcal{R}_k . Then, for any convex object C in the plane containing at least εn points, either a) $C \cap Q \neq \emptyset$, or b) there exist $k/2$ points, say $q_i \in Q, i = 1 \dots k/2$, such that C intersects all edges $\overline{q_i q_j}$ for all $1 \leq i < j - 1 \leq k$.*

Proof. Assume $C \cap Q = \emptyset$. We then give a deterministic procedure that always finds k such points. W.l.o.g. assume that the convex object is polygonal, and denote its vertices in cyclic order by

p_1, \dots, p_m for some m . Note that the next vertex after p_m is p_1 again.

Define $\overrightarrow{p_i p_{i+1}}$ as the (infinite) half-line starting at p_i , containing p_{i+1} and extending towards p_{i+2} (define $\overrightarrow{p_{i+1} p_i}$ likewise). See Figure 3.1 (a). Let $T(i, j)$ be the region bounded by $\overrightarrow{p_{i-1} p_i}$, the segments $p_i p_{i+1}, \dots, p_{j-1} p_j$, and $\overrightarrow{p_{j+1} p_j}$. Initially set $l = 1$, $i_l = 2$, and $j = 3$, and repeat the following:

1. If $T(i_l, j)$ contains a point of Q , denote this point (pick an arbitrary one if there are many) to be q_l . Set $i_{l+1} = j$. Increment l to $l + 1$, set $j = j + 1$, and continue as before to find the next point of Q .
2. If $T(i_l, j)$ does not contain any point of Q , extend the region by incrementing j to $j + 1$, and check again if $T(i, j)$ contains a point of Q .

This process ends when $j = 1$. Assume we have l points q_1, \dots, q_l , together with the indices i_1, \dots, i_l . Note that, by construction, each point q_t is contained in the region $T(i_t, i_{t+1})$. Consider any i_t and the point q_t that the region $T(i_t, i_{t+1})$ contains. See Figure 3.1(b).

Claim 3.2. *The region $T(i_{t-1}, i_t - 1)$ has no points of Q .*

Proof. By the greedy method of construction, i_t is the smallest index j for which the region $T(i_{t-1}, j)$ is non-empty. Hence all the regions $T(i_{t-1}, j)$, $i_{t-1} < j < i_t$ are empty. \square

Define h_t to be the halfspace incident to the edge $p_{i_t-1} p_{i_t}$ and containing \mathcal{C} . Claim 3.2 immediately implies the following.

Claim 3.3. *The halfspace h_t , defined by the line incident to the edge $p_{i_t-1} p_{i_t}$, separates q_t (and all the other points of Q) lying in $T(i_{t-1}, i_t)$ from \mathcal{C} .*

If the number of points found by our method is at most k (i.e., $l \leq k$), then take the intersection of the half-spaces h_t , for $t = 1 \dots l$. By Claim 3.3, each halfspace h_t separates all the points in $T(i_{t-1}, i_t)$ from \mathcal{C} . Thus all the points of Q are now separated by this intersection (see Figure 3.1 (a) for the separating halfspaces), and since each halfspace contains \mathcal{C} , the intersection contains at least εn points of P . This is a contradiction to the fact that Q was a ε -net to the range space \mathcal{R}_k .

Finally, note that the sequence q_t of points obtained, $t = 1 \dots k$, has the property that the intersection point of any (properly intersecting) pair of segments joining non-consecutive points, lies inside \mathcal{C} . This follows from the fact that for every point q_t , all the non-adjacent points and q_t lie in the same two half-spaces incident to edges $p_{i_t-1} p_{i_t}$ and $p_{i_{t+1}} p_{i_{t+1}+1}$, both of which are incident to \mathcal{C} . Therefore picking every alternate point yields the desired set. \square

Set $k = 8$, and compute the ε -net for the range space \mathcal{R}_8 . By Lemma 3.1, there exists a sequence of four points, say a, b, c, d , such that \mathcal{C} contains the intersection of the two segments ac and bd . This immediately yields a way to construct weak ε -nets using (strong) ε -nets: compute a ε -net for \mathcal{R}_8 , and add the intersection points of all segments between pairs of points. By the above argument, each convex object contains at least one of these points. The number of points in the weak ε -net constructed above are $O(1/\varepsilon^4 \log^4 1/\varepsilon)$. We now show that by a more careful argument, this can be reduced to $O(1/\varepsilon^3 \log^3 1/\varepsilon)$.




Figure 3.1: Constructing weak ε -nets in two dimensions. (a) The dotted red lines indicate the at most k halfspaces that are used to separate Q from C .

Theorem 3.4. Given a set P of n points in the plane, construct an ε -net Q for the range space $\mathcal{R}_{1,2}$. Construct the set Q' as follows: for every ordered triple of points in Q , say a, b, c , add the intersection of the bisector of $\angle abc$ to the line segment ac . Then Q' has size $O(1/\varepsilon^3 \log^3 1/\varepsilon)$ and is a weak ε -net for P .

Proof. Fix a convex object C containing at least εn points of P . From Lemma 3.1, there exists a sequence of six points in convex position, say a, b, c, d, e, f , of Q where the intersection point of every pair of (properly intersecting) segments spanning these points lies in C .

The sum of the interior angles of the polygon defined by the six points is 4π . Form two triangles by taking alternate points, say $\triangle ace$ and $\triangle bdf$. The sum of the interior angles of the two triangles is 2π . By pigeon-hole principle, there exists a point, say a , where the angle $\angle cae$ is at least one-half of the interior angle of the polygon at vertex a , $\angle fab$. Therefore, the bisector of the interior angle $\angle fab$ lies inside the triangle ace , and intersects the segment bf . This intersection lies between the intersection of bf with the two segments ac and ae . See Figure 3.2(a). By assumption, these two intersections are contained inside C . Therefore, by convexity, the intersection of the bisector of $\angle fab$ with the segment fb lies inside C . Since Q' contains all such intersections, C is hit by Q' . \square

3.2 Three Dimensions

Lemma 3.5. Given a positive integer t , a convex set C and a set of points Q in \mathbb{R}^d such that $C \cap Q = \emptyset$, then i) either the set Q can be separated from C by $f(t)$ hyperplanes or ii) there exists $Q' \subseteq Q$ such that $|Q'| = t$ and the convex hull of every $(d+1)$ points of Q' intersects C .

Proof. For any point $q \in Q$, define the set

$$S(q) = \{\vec{a} \in \mathbb{R}^d \mid \vec{a} \cdot \vec{q} \geq \vec{a} \cdot \vec{x}, \forall x \in \mathcal{C}\}$$

First note that $S(q)$ is a convex set: take any two vectors $\vec{a}, \vec{b} \in S(q)$. Then,

$$\begin{aligned} \vec{a} \cdot \vec{q} &\geq \vec{a} \cdot \vec{x}, \forall x \in \mathcal{C} \\ \vec{b} \cdot \vec{q} &\geq \vec{b} \cdot \vec{x}, \forall x \in \mathcal{C} \\ \Rightarrow (\lambda \vec{a} + (1 - \lambda) \vec{b}) \cdot \vec{q} &\geq (\lambda \vec{a} + (1 - \lambda) \vec{b}) \cdot \vec{x}, \forall x \in \mathcal{C} \forall \lambda \in [0, 1], \end{aligned}$$

implying that $S(q)$ is a convex object in \mathbb{R}^d .

By definition, $a \in S(\vec{q})$ implies that there is a hyperplane whose normal is parallel to \vec{a} and which separates q from the \mathcal{C} . If there are $d + 1$ points q_1, \dots, q_{d+1} whose convex hull does not intersect \mathcal{C} , then these $d + 1$ points can be separated from \mathcal{C} by a single hyperplane. This implies that the corresponding convex objects $S(q_1), \dots, S(q_{d+1})$ have a common intersection.

Let $S = \cup_{q \in Q} S(q)$ be the set of convex objects corresponding to the points in Q . If every subset $Q' \subseteq Q$ of size t has $(d + 1)$ points whose convex hull does not intersect \mathcal{C} , then $(d + 1)$ of every t convex objects in S intersect. Therefore applying the (p, q) -Hadwiger Debrunner theorem with $p = t$ and $q = (d + 1)$ on the convex sets in S , we deduce that Q can be separated from \mathcal{C} using $f(t)$ hyperplanes, where $f(t) = HD_d(t, d + 1)$ and $HD_d(p, q)$ is the Hadwiger-Debrunner hitting set number for p and q in d dimensions. \square

Remark 1: One can also see the statement in the dual setting. For a point $q \in Q$, the dual of the space of hyperplanes separating q from \mathcal{C} is a convex object. Then a hyperplane separating t points of Q from \mathcal{C} becomes a point in the dual space which is in the common intersection of the t dual convex objects for these points.

Remark 2: We will prove a stronger statement in the next chapter which will imply that the convex-hull of every $\lfloor d/2 \rfloor + 1$ tuple of Q' will intersect \mathcal{C} .

Lemma 3.6. *Given a convex object \mathcal{C} and a set Q' of $g(t)$ points in \mathbb{R}^3 , such that the convex hull of every 4 points in Q' intersects \mathcal{C} , we can find $Q'' \subseteq Q'$ of size at least $t \geq 5$ such that the convex hull of every 3 points in Q'' intersects \mathcal{C} .*

Proof. Consider a hypergraph with the base set Q' and every 3-tuple of points in Q' as a hyperedge. Color a hyperedge ‘red’ if the convex hull of the corresponding 3 points intersects \mathcal{C} and ‘blue’ otherwise. Then, by Ramsey’s theorem for hypergraphs [Die00], there exists a constant $g(t)$ such that if $|Q'| \geq g(t)$, there exists a monochromatic clique, say Q'' , of size t . A monochromatic ‘blue’ clique implies that there exists a set of t points such that \mathcal{C} does not intersect the convex hull of any 3-tuple of these points. Take any 5 points of Q'' , and partition their convex hull into two tetrahedra sharing a face. Since both these tetrahedra must intersect \mathcal{C} , their common face must also intersect \mathcal{C} , a contradiction. Therefore, the clique returned must be monochromatic ‘red’, implying the existence of a subset Q'' of size t such that the convex hull of all three points in Q'' intersects \mathcal{C} . \square




Figure 3.2: (a) The intersection of a bisector with a segment will lie inside \mathcal{C} , (b) If \mathcal{C} intersects edges ac , ad and ae , then it must intersect af . Similarly for bf .

To prepare for the next lemma, we need the following geometric claim.

Claim 3.7. *Let $T = \{a, b, c, d, e\}$ be a set of five points in convex position in \mathbb{R}^3 . Then, if a convex object \mathcal{C} intersects the convex hull of every 3-tuple of T , it intersects at least one edge (convex hull of a 2-tuple) spanned by the points in T .*

Proof. By Radon's theorem, in every set of five points in convex position, there exists a line segment which intersects the convex hull of the remaining three points (the Radon partition). Assume the line segment ab intersects the convex hull of c, d , and e . Then, we claim that \mathcal{C} must intersect ab . Otherwise, there exists a hyperplane h separating ab from \mathcal{C} . Since ab intersects the convex hull of c, d and e , h separates at least one point in $\{c, d, e\}$ from \mathcal{C} and convex hull of a, b and this third point does not intersect \mathcal{C} , a contradiction. \square

Lemma 3.8. *Given a convex objects \mathcal{C} and a set Q'' of $h(t)$ points such that the convex hull of every 3 points in Q'' intersects \mathcal{C} , there exists a subset $Q''' \subseteq Q''$ of size t such that the convex hull of every two points in Q''' intersects \mathcal{C} .*

Proof. Again consider a hypergraph with the base set Q'' and every 2-tuples of these points as a hyperedge. Color a hyperedge 'red' if the convex hull of the corresponding 2-tuple intersects \mathcal{C} and 'blue' otherwise. Then again by Ramsey's theorem, there exists a positive integer $h(t)$ such that if $|Q''| \geq h(t)$, there exists a monochromatic clique of size t . We can assume (again by Ramsey's theorem) that if $t \geq k$ where k is a constant, then the points of the monochromatic clique have 5 points in convex position. From Claim 3.7, it follows that the convex hull of two of the points of these 5 points intersects \mathcal{C} , thereby implying that the color of the monochromatic clique cannot be 'blue' and hence the convex hull of every pair of points in the clique intersects \mathcal{C} . \square

Lemma 3.9. *Given a set of points R in convex position, $|R| \geq 5$, and a convex object \mathcal{C} that intersects every edge spanned by the points in R , the Radon point of R is contained in \mathcal{C} .*

Proof. Take the Radon partition of any five points in R . See Figure 3.2 (b). Say the edge ab intersects the facet spanned by $\{c, d, e\}$. It is easy to see that if C intersects the edges ac , ad and ae , it must intersect the segment af . Similarly, if C intersects the edges bc , bd and be , it intersects the segment bf . By convexity, it must contain the intersection of the edge ab with $\triangle cde$. \square

We come to our main theorem in this section:

Theorem 3.10. *Let P be a set of n points in \mathbb{R}^3 . Then there exists a constant c (function of the dimension, 3 in this case) such that the followings holds: take any ε -net, say Q , with respect to the range space (P, \mathcal{R}_c) . Construct a weak ε -net, say Q' , as follows: for every ordered 5-tuple, say a, b, c, d, e , add the intersection (if any) of $\triangle abc$ with \overline{de} . Then Q' is a weak ε -net for P of size $O(1/\varepsilon^5 \log^5 1/\varepsilon)$.*

Proof. Fix any convex set C containing at least εn points of P . For a large enough constant c (depending on $f(\cdot), g(\cdot), h(\cdot)$), by Lemma 3.5, Lemma 3.6 and Lemma 3.8, there exists a set of at least five points such that C intersects every edge spanned by these points. Lemma 3.9 then implies that Q' is a weak ε -net. \square

Remark: In [MW04], in order to construct a set that contains a centerpoint of all subsets of a set of r points in d dimensions, r^{d^2} points are used. The techniques described above can be used to reduce this to r^3 and r^5 (instead of r^4 and r^9) for dimensions two and three respectively. This improves the logarithmic factors in their result.

3.3 Higher Dimensions

Although the optimal weak ε -net can consist of any subset of \mathbb{R}^d , arguing similar to [MW04], we show that there is a discrete finite set of points in \mathbb{R}^d from which an optimal weak ε -net can be chosen. Given P , this subset is constructed as follows: consider the set of all hyperplanes spanned by the points of P (each such hyperplane is defined by d points of P). Every d of these hyperplanes intersect in a point in \mathbb{R}^d , and consider all such points formed by intersection of every d hyperplanes. This is the required point set, which we denote by $\Xi(P)$.

Lemma 3.11. *Let P be a set of n points in \mathbb{R}^d . Then the set $\Xi(P)$, of size $O(n^{d^2})$, contains an optimal weak ε -net for P , for any $\varepsilon > 0$.*

Proof. Let S be any weak ε -net for P . We show how to locally move each point of S to a point of $\Xi(P)$. Wlog assume that each convex set is the convex hull of the points it contains. Take a point $r \in S$, and consider the (non-empty) intersection of all the convex sets which contain r . The lexicographically minimum point of this intersection, t , is the intersection of d of these convex objects [Mat02]. Note that t lies on a facet of each of these convex objects, and each facet is a hyperplane passing through d points of P . Replacing r with t still results in a weak net, since by construction, t is also contained in all the convex objects containing r . The proof follows. \square

We now show that $\Xi(Q)$, where Q is a random sample of P of size $O(1/\varepsilon \log 1/\varepsilon)$, is a weak ε -net with constant probability.

Theorem 3.12. *Let P be a set of n points in \mathbb{R}^d , and let Q be a random sample of size $O(1/\varepsilon \log 1/\varepsilon)$ from P . With constant probability, $Q' = Q \cup \Xi(Q)$ is a weak ε -net for P .*

Proof. Clearly Q' has size $O(\varepsilon^{-d^2} \log^{d^2} 1/\varepsilon)$ since each point in Q' is defined by at most d^2 points of Q (intersection of d hyperplanes, each defined by d points).

First, with constant probability, Q is an ε -net with respect to the range space (P, \mathcal{R}_c) , where $c = f((d+1)^2)$. Let \mathcal{C} be any convex object containing at least εn points of P and assume $\mathcal{C} \cap Q = \emptyset$. Then \mathcal{C} cannot be separated from Q by c hyperplanes, otherwise the intersection of the halfspaces containing \mathcal{C} defined by these c hyperplanes has εn points and no point of Q , a contradiction to the fact that Q is an ε -net for (P, \mathcal{R}_c) . By Lemma 3.5, there exist a set S of at least $(d+1)^2$ points of Q such that the convex hull of every $(d+1)$ of them intersects \mathcal{C} .

By Lemma 1 of [MW04], Q' contains a centerpoint, say q , of the set S . We claim that q is contained in \mathcal{C} . Otherwise, by the separation property, there exists a halfspace h^- containing q such that $h^- \cap \mathcal{C} = \emptyset$. By the centerpoint property, h^- contains at least $(d+1)^2/(d+1) = (d+1)$ points of S . The convex hull of these $(d+1)$ points lies in h^- and therefore does not intersect \mathcal{C} , a contradiction. \square

Given a set Q , a *deep-point* is a point $q \in \mathbb{R}^d$ such that any hyperplane containing q contains at least d points of Q . Let $c(Q)$ be the set of points in \mathbb{R}^d such that a deep-point of every subset of Q is present in $c(Q)$. The proof above implies the following.

Corollary 3.3.1. *If $c(Q)$ has size $O(m^t)$ for any set Q , one can construct a weak ε -nets for any pointset of size $\tilde{O}(1/\varepsilon^t)$.*

This page was intentionally left blank

4. A Generalization of Carathéodory's Theorem

Counting pairs is the oldest trick in combinatorics . . . every time we count pairs, we learn something from it.

Gil Kalai

In this chapter we present lower-dimensional extensions of three basic theorems: Radon's theorem, Carathéodory's theorem, and Colorful Carathéodory's theorem.

The proof of low-dimensional extension of Radon's Theorem follows trivially from this well-known generalization of the Erdős-Szekeres theorem:

Theorem 4.1 (Generalized Erdős-Szekeres Theorem). *Given positive integers d, k, n such that $\lceil d/2 \rceil + 1 \leq k \leq d$, there exists an integer $n_0 = ES_d(n, k)$ such that any set of n_0 points in \mathbb{R}^d contains a subset P of size n with the following property: the simplex spanned by every $(d + 1) - k$ points of P lies on the boundary of $\text{conv}(P)$. This statement is optimal, in the sense that this is not true for $k < \lceil d/2 \rceil + 1$ for arbitrarily large pointsets.*

The case $k = d$ simply corresponds to the Erdős-Szekeres theorem (that any large-enough set contains a lot of points in convex position). Of course the 'large-enough' size for the above theorem increases with decreasing k ; but if one pays that price, one can get more properties. For example, for $d = 4, k = 3$, any large-enough set of points in \mathbb{R}^4 contains a large subset Q where every edge spanned by points of Q lies on $\text{conv}(Q)$.

We now observe that this immediately carries over to an extension of Radon's theorem: if one is willing to increase the number of points, then a better upper-bound can be achieved on the sizes of the Radon partition:

Theorem 4.2. *Given an integer $\lfloor d/2 \rfloor + 1 \leq k \leq d$, any set P of $ES_d(d + 2, k)$ points in \mathbb{R}^d contains two sets P_1, P_2 such that $\text{conv}(P_1) \cap \text{conv}(P_2) \neq \emptyset$ and additionally, $|P_1|, |P_2| \leq k$. Furthermore, this is optimal in the sense that the statement does not hold for $k \leq \lfloor d/2 \rfloor$.*

Proof. Apply Theorem 4.1 to P to get a set of $d + 2$ points P' . Apply Radon's theorem to P' to get a partition $P_1, P_2 \subset P'$ whose convex hulls intersect. Now note that if $|P_1| > k$, then $|P_2| \leq (d + 1) - k$. But then $\text{conv}(P_2)$ lies on the convex hull of P' , and so cannot intersect $\text{conv}(P_1)$, a contradiction.

Optimality is obvious as $|P| \geq d + 2$ for such a partition to exist (for P in general position), and so one set has to have at least $\lfloor d/2 \rfloor + 1$ points. \square

4.1 Proof of Extended Carathéodory's Theorem

The goal of this section is to prove the following:

Theorem 4.3. *Given a set P of n points in \mathbb{R}^d and a convex object C , either P can be separated from C by $O(d^4 \log d)$ hyperplanes (i.e., each $p \in P$ is separated from C by one of the hyperplanes), or C intersects the convex hull of some $\lfloor d/2 \rfloor + 1$ -sized subset of P .*

We first show that this problem is related to another problem involving low-dimensional simplices.

Let $f(d)$ be the smallest positive number such that for any set P of points in \mathbb{R}^d , there exists a $\lfloor d/2 \rfloor + 1$ -sized subset $P' \subseteq P$ such that any halfspace containing P' contains at least $|P|/f(d)$ points of P .

Let $g(d)$ be the smallest positive number such that given any set P of points in \mathbb{R}^d and a convex set C , if P cannot be separated from C using at most $g(d)$ hyperplanes, then C must intersect the convex hull of some $\lfloor d/2 \rfloor + 1$ size subset of P .

We now show that $g(d)$ and $f(d)$ are related within a factor of d .

Theorem 4.4. $g(d) \leq d \cdot f(d) \log f(d)$. *In other words, given a set P of points in \mathbb{R}^d and a convex set C such that P cannot be separated from C by $df(d) \log f(d)$ hyperplanes, then C must intersect the convex hull of some $\lfloor d/2 \rfloor + 1$ points of P .*

Proof. Assume that no convex hull of any $\lfloor d/2 \rfloor + 1$ points of P intersects C . Then we show that P can be separated from C using $df(d) \log f(d)$ hyperplanes.

Claim 4.5. *Let P be a weighted set of points in \mathbb{R}^d , with weight of the point $p_i \in P$ to be w_i . Assume all w_i 's are rationals, and let $W = \sum w_i$. If the convex hull of no $\lfloor d/2 \rfloor + 1$ points of P intersects C , then there exists a hyperplane separating points of total weight at least $\frac{W}{f(d)}$ from C .*

Proof. As each w_i is a rational, assume $w_i = \hat{w}_i/D$, where \hat{w}_i and D are integers. Let Q be the pointset gotten by replacing each point p_i with \hat{w}_i copies of p_i . Crucially, if the convex hull of no $\lfloor d/2 \rfloor + 1$ subset of P intersects C , then the convex hull of no $\lfloor d/2 \rfloor + 1$ subset of Q can intersect C . Take the $\lfloor d/2 \rfloor + 1$ -sized subset Q' of Q such that any halfspace containing Q' contains at least $|Q|/f(d)$ points of Q . As the convex hull of Q' does not intersect C , there is a halfspace h which does not intersect C and contains Q' . Let P' be the set of points of P contained in h . Then h contains exactly $\sum_{p_i \in P'} \hat{w}_i$ points of Q , which by definition of Q' must be at least $|Q|/f(d)$. Then the sum of weights of points of P contained in h is bounded by

$$\begin{aligned} \sum_{p_i \in P'} w_i &= \frac{\sum_{p_i \in P'} \hat{w}_i}{D} \geq \frac{|Q|/f(d)}{D} = \frac{(\sum \hat{w}_i)/f(d)}{D} \\ &= \frac{\sum w_i D}{D f(d)} = \frac{W}{f(d)} \end{aligned}$$

□

Discretize the set of all combinatorially distinct hyperplanes separating some subset of P from C to get a set $\mathcal{H} = \{h_1, \dots, h_m\}$ of $O(|P|^d)$ hyperplanes. Now consider assigning weights $w(h_i)$ to each halfspace such that the total weight $\sum w(h_i)$ is minimized, and the sum of weights of halfspaces containing any point of P is at least 1. Let $W(\mathcal{H})$ denote the minimum value.

Similarly, assign weights $w(p_i)$ to each point of P such that the total weight $\sum w(p_i)$ is maximized, and the sum of weights of points contained in any halfspace $h \in \mathcal{H}$ is at most 1. Let $W(P)$ denote the maximum value. Then the above two problems are dual to each other (as linear programs), and so by the Strong Duality Theorem, $W(\mathcal{H}) = W(P)$.

Now note that $W(P) \leq f(d)$: by Claim 4.5, there exists a halfspace in \mathcal{H} of weight at least $W(P)/f(d)$, which by the definition of $W(P)$ is at most 1.

Therefore there exists an assignment of weights to halfspaces in \mathcal{H} such that $W(\mathcal{H}) \leq f(d)$, and each point is contained in halfspaces of total weight at least 1. Now using the ε -net theorem for halfspaces [KPW92], with $\varepsilon = 1/W(\mathcal{H})$, there exists a set of

$$d/\varepsilon \log 1/\varepsilon = dW(\mathcal{H}) \log W(\mathcal{H}) = df(d) \log f(d)$$

halfspaces of \mathcal{H} containing all points of P . As all halfspaces in \mathcal{H} were separating halfspaces, we are done. □

Remark 1: The above technique is similar to the one used in the proof of Hadwiger-Debrunner (p, q) theorem [AK92], with some crucial differences. In their use, they get an exponential bound, which we are able to avoid due to three reasons: ε -nets for halfspaces have a near-linear bound, avoiding double-counting arguments that they use, and finally, the weighted version (Claim 4.5) gives exactly the same quantitative bound as the unweighted version.

Proof of Theorem 4.3: the paper [SSW08] proves that $f(d) \leq O(d^3)$. And the proof is complete by using Theorem 4.4.

Similarly we now show that a bound on $g(d)$ gives an upper-bound on $f(d)$. We will need the following fact:

Fact 1 [PA95]: If P is a set of n points and h is a hyperplane defining a facet of the β -deep region C of P , then the halfspace defined by h that does not intersect the interior of C contains less than βn points of P .

Theorem 4.6. $f(d) \leq d \cdot g(d)$. In other words, given a set P of points in \mathbb{R}^d , there always exists a subset P' of size $\lfloor d/2 \rfloor + 1$ such that any halfspace containing P' contains at least $|P|/dg(d)$ points of P .

Proof. Consider the β -deep region C of P ; by the Centerpoint theorem, for $\beta \leq 1/(d+1)$, such a region always exists. Now we claim that for $\beta = 1/dg(d)$, there exists a $\lfloor d/2 \rfloor + 1$ -sized subset P' whose convex hull intersects C . Then any halfspace containing P' contains at least one point of C , and so contains at least $|P|/dg(d)$ points by the definition of the centerpoint region.

Otherwise, for contradiction assume that the convex hull of no $\lfloor d/2 \rfloor + 1$ -sized subset intersects C . Then by definition of $g(d)$, P can be separated from C using $g(d)$ hyperplanes, say the set \mathcal{H} .

Now any halfspace not intersecting C contains less than $d \cdot \beta|P|$ points: each halfspace supporting a facet of C contains less than $\beta|P|$ points, and any other halfspace not intersecting C is contained in the union of at most d halfspaces supported by facets of C .

Therefore each halfspace of \mathcal{H} contains less than $d \cdot \beta|P|$ points. And so the union of halfspaces in \mathcal{H} contains less than $g(d) \cdot d \cdot \beta|P|$ points of P , a contradiction for $\beta = 1/dg(d)$. \square

4.2 Proof of Extended colorful Carathéodory Theorem

The goal of this section is to prove:

Theorem 4.7. *For any d , there exists a constant N_d such that given $k = \lfloor d/2 \rfloor + 1$ sets of points P_1, \dots, P_k in \mathbb{R}^d and a convex object C , either one of the sets P_i can be separated from C by N_d hyperplanes, or there is a rainbow set of size k whose convex hull intersects C .*

The approach of the previous section can be made to work for proving the extension of the colorful Carathéodory's theorem through the use of Ramsey-theoretic techniques. Alternatively, we now present a different proof which highlights explicitly the connection to the Hadwiger-Debrunner theorem of Alon-Kleitman [AK92].

We use a slightly different language for convenience: instead of saying that “a point set P can be separated from a convex body C using k hyperplanes”, we say that “there exists a polyhedron Q with k facets such that $C \subseteq Q$ and $Q \cap P = \emptyset$ ”. In such a case we also say that Q separates P from C . We re-state Theorem 4.7 in this language.

Theorem 4.8. *For any positive d and $l > \lfloor d/2 \rfloor$, there exists a constant $N_{d,l}$ s.t. the following is true. Given any compact convex body C and l finite sets of points P_1, \dots, P_l in \mathbb{R}^d , at least one of the following holds:*

1. *There exists a polyhedron Q with at most $N_{d,l}$ facets such that for some i , Q separates P_i from C .*
2. *There exists a rainbow subset $P' \subseteq \cup_{i=1}^d P_i$ whose convex hull intersects the interior of C .*

Call a convex body C *fine* if it is compact and its boundary ∂C is smooth and has positive curvature everywhere. Let C be a fine convex body and let P be a finite set of points in \mathbb{R}^d . We say that a point p can see a point y if the relative interior of the segment py does not intersect C . For any $p \in P$, let U_p be the set of points in ∂C that p can see.

Let h_y be the tangent plane to C at the point $y \in \partial C$ and let h_y^+ be the closed halfspace defined by it that contains C . Observe that any point $p \in P$ sees a point $y \in \partial C$ iff $p \notin \text{int}(h_y^+)$, where $\text{int}(S)$ denotes the interior of the set S .

Lemma 4.9. *For any positive numbers d and $t \geq d$, there exists a constant $H_{t,d}$ such that given any fine convex body C and a finite set of points P such that $P \cap C = \emptyset$, at least one of the following hold:*

- There exists a set $X \subset \partial C$ of size at most $H_{t,d}$ such that each point in P is seen by some $x \in X$.
- There is a subset $P' \subseteq P$ of size at least t such that no $y \in \partial C$ sees more than $d - 1$ points in P' .

Proof. Fix any point $\nu \in \partial C$ as a reference point. Let h_ν be the tangent plane to C at ν . Let h be the unique tangent hyperplane parallel to h_ν such that C is contained in the strip between h_ν and h . Let \prec be the continuous bijective map that maps any $y \in \partial C \setminus \{\nu\}$ to $l(\nu, y) \cap h$, where $l(\nu, y)$ denotes the line through ν and y .

For any $p \in P$ that does not see ν (i.e., $\nu \notin U_p$), let $V_p = \prec(U_p) = \{\prec(y) : y \in U_p\}$ and let $\mathcal{V} = \{V_p : p \in P, p \text{ does not see } \nu\}$. Clearly, for each $p \in P$ that does not see ν , V_p lies on the plane h , is convex and has dimension $d - 1$. Note that if two points $p, q \in P$ see a point $y \in \partial C$ then $y \in U_p \cap U_q$. So $\prec(y) \in V_p \cap V_q$.

Suppose that the second part of the theorem does not hold; i.e., in every subset of P of size t , there are at least d points which can be seen by a single point $y \in \partial C$. Equivalently, any subset of \mathcal{V} of size t has at least d sets which have a common intersection. By the Hadwiger-Debrunner theorem [AK92], there exists a constant $HD_{d-1}(t, d)$ of points in h that hit all the sets in \mathcal{V} . Let X' be the set of these points. Let $X = \prec^{-1}(X') \cup \{\nu\}$. Each point in P is seen by at least one point in X (if V_p is hit by the point $y' \in h$, then $p \in P$ is seen by $\prec^{-1}(y') \in \partial C$). The theorem is therefore proved by setting $H_{t,d} = HD_{d-1}(t, d) + 1$. \square

Now we can finish the proof of the main theorem of this section:

Proof. Let $P = \cup_{i=1}^l P_i$. We set $N_{d,l} = H_{t,d}$ for some t to be fixed later.

If $P \cap C \neq \emptyset$ then the second part of the theorem is trivially satisfied. We therefore assume that $P \cap C = \emptyset$. Without loss of generality we also assume that C is fine since we can always find a fine convex body C' that contains C and does not intersect P and furthermore for each point $y' \in C'$, there is a point $y \in C$ such that the Euclidean distance between y and y' is smaller than any prescribed $\delta > 0$. Proving the theorem for such C' 's also proves it for arbitrary closed convex bodies.

For each i , apply Lemma 4.9 to C and P_i with the parameter t . This gives us either a set X_i of at most $H_{t,d}$ points in ∂C such that each point in P_i is seen by at least one of these or we get a set $Q_i \subseteq P_i$ of t points so that no d of them is seen by the same point in ∂C . If the first possibility happens for some j , then $\bigcap_{x \in X_j} h_x^+$ gives us the polyhedron Q with at most $H_{t,d}$ facets and where Q contains C while P_j lies outside Q . This satisfies the first part of the Theorem and we're done.

We therefore assume the second possibility for each i ; namely, each P_i has a subset Q_i of t points such that no d points of Q_i are seen by the same point of ∂C . Equivalently, the convex hull of any d points of Q_i intersects C .

Let $Q = \cup_{i=1}^l Q_i$. Consider any rainbow set $R \subseteq Q$ with one point from each Q_i . There are t^l such sets. If the convex hull of R intersects C , then the second part of the theorem is satisfied, and we're done. Assume for contradiction that this is not the case. Then for each rainbow set R , there exists a hyperplane h separating R from C . The closed halfspace h^- bounded by h and not intersecting C contains at most $d - 1$ points from any particular Q_i due to the fact that any d -sized subset of Q_i intersects C . Therefore $|h^- \cap Q| \leq (d - 1)l$, and hence is a $\leq k$ -set of Q with $k = l(d - 1)$.

If no rainbow set intersects C , then we get such a $\leq k$ -set for each rainbow set R of size l . As there are t^l such rainbow sets, we get $t^l \leq k$ -sets. However each such $\leq k$ -set can be overcounted at most $\binom{k}{l} = \binom{l(d-1)}{l}$ times. This implies that there are at least $L(t) = t^l / \binom{l(d-1)}{l}$ distinct $\leq k$ -sets. On the other hand, it is known that the number of $\leq k$ -sets of a set of n points in \mathcal{R}^d is at most $O(n^{\lfloor d/2 \rfloor} (k+1)^{\lceil d/2 \rceil})$ [Mat02, p. 265]. This gives an upper bound of $U(t) = O((tl)^{\lfloor d/2 \rfloor} ((d-1)l+1)^{\lceil d/2 \rceil})$ on the number of $\leq k$ -sets. Since $l > \lfloor d/2 \rfloor$, for some large enough t depending only on l and d , $L(t) > U(t)$. Thus we get a contradiction implying that one of the rainbow sets must intersect C . \square

Part II

Algorithms

5. Independent sets in Intersection Graphs of Rectangles

I want to express a radical alternative that I learned from Sir Michael Atiyah. His view was that the most significant aspects of a new idea are often not contained in the deepest or most general theorem which they lead to. Instead, they are often embodied in the simplest examples, the simplest definition and their first consequences.

David Mumford

Let $S = \{s_1, \dots, s_n\}$ be a set of n rectangles in \mathbb{R}^2 , and let $\alpha(S)$ denote the size of the largest independent set in the intersection graph of S . Similarly, let $\omega(S)$ denote the size of the largest clique in the intersection graph of S .

In this chapter we look at the independent-set problem in the intersection graph of axis-parallel rectangles in the plane, and prove the following main theorem.

Theorem 5.1. *Given S with $\alpha(S) = \beta n$, one can compute an independent set of size $\frac{\alpha(S)}{4d_0 \cdot (1/\beta)}$ in polynomial time, where d_0 is a constant.*

Without loss of generality, we can assume that no rectangle completely contains any other rectangle. We first consider two special cases of rectangle intersection graphs — the *piercing* and *non-piercing* intersection subgraphs, derived by defining the following partial order among the rectangles.

Given two rectangles s_1 and s_2 , $s_1 \prec s_2$ if and only if s_1 intersects both vertical edges of s_2 , and s_2 intersects both horizontal edges of s_1 (See Figure 5.1(a)). Clearly, if $s_1 \prec s_2$, then s_1 and s_2 intersect, and neither contains any vertex of the other. Note that \prec is a transitive relation: $s_1 \prec s_2 \prec s_3$ implies $s_1 \prec s_3$. If $s_1 \prec s_2$, we say that s_1 and s_2 *pierce*, and that s_2 is pierced by s_1 .

Piercing and Non-Piercing Rectangle Intersection Graphs

Given S , the relation \prec partitions the edges in the intersection graph G_S into two sets E_1 and E_2 : given an edge $\{a, b\} \in E(G_S)$, $\{a, b\} \in E_1$ if $a \prec b$ or $b \prec a$, and $\{a, b\} \in E_2$ otherwise.




Figure 5.1: (a) $s_1 \prec s_2 \prec s_3$, (b) The set S' (solid) and the set I' (dashed), (c) Mapping rectangles in S' to vertices, and rectangles in I' to edges.

Piercing Intersection Graphs. Define the piercing intersection graph of S as the directed graph $G_p = (V, E)$, i.e., each vertex in $V(G_p)$ corresponds to a rectangle in S , and there is a directed edge (a, b) between two vertices a and b if b is pierced by a , i.e., $(a, b) \in E(G_p)$ if and only if $a \prec b$.

Lemma 5.2. Given a set S of n rectangles in \mathbb{R}^2 , let G_p be the piercing graph of S . Then $\omega(G_p) \cdot \alpha(G_p) \geq n$, and a maximum independent set in G_p can be computed in polynomial time.

Proof. It follows from the transitivity of \prec that G_p is a transitive graph. It is a well-known fact that all transitive graphs are *perfect* graphs [Gol04], i.e., for every induced subgraph of G_p , the size of the maximum clique equals the chromatic number. Therefore the vertices of G_p can be partitioned into $\omega(G_p)$ subsets, each of which is an independent set in G_p . Since the largest subset has at least $n/\omega(G_p)$ vertices, we conclude that $\omega(G_p) \cdot \alpha(G_p) \geq n$.

By a classical result of Grötschel *et al.* [GLS81], a largest independent set of a perfect graph can be computed in polynomial time. \square

Non-Piercing Intersection Graphs. We now consider the case of pairwise non-piercing rectangles. Let S be a set of n rectangles such that no two rectangles pierce (although they could intersect); the intersection graph of S is called the non-piercing intersection graph and denoted as G_{np} . First, note that computing the optimal independent set of S remains NP-hard since the construction in [RN95] proving the NP-hardness for intersection graphs of rectangles uses only non-piercing rectangles.

We call a subset $S' = \{s_{i_1}, \dots, s_{i_m}\} \subseteq S$ *r-maximal* if it satisfies the following four conditions:

- (A1) S' is an independent set,
- (A2) Every $s \in S \setminus S'$ intersects some $s' \in S'$,
- (A3) For each $s' \in S'$, there is at most one rectangle $s \in \text{OPT}(S)$ such that s intersects s' and does not intersect any other rectangle in S' , and

(A4) For every pair of disjoint rectangles $s', t' \in S'$, there are at most two mutually-disjoint rectangles $s, t \in S \setminus S'$ such that s, t intersect both s' and t' , and no other rectangle in S' .

Lemma 5.3. *Let $S' = \{s_{i_1}, \dots, s_{i_m}\} \subseteq S$ be an r -maximal set. Then $|S'| \geq \alpha(S)/c_0$, where $c_0 \leq 11$.*

Proof. Let $I = \text{OPT}(S)$ be the set of rectangles in a maximum independent set. We will charge each rectangle in I to a rectangle in S' such that each rectangle in S' is charged at most a constant c_0 times. The lemma then follows.

Let s be a rectangle in I . By maximality condition (A2), s intersects at least one rectangle in S' , and since the rectangles in S are pairwise non-piercing, s does not pierce any rectangle in S' . We charge s to a rectangle in S' as follows. First, if $s \in S'$, we charge s to itself. Clearly, each rectangle in $S' \cap I$ receives only one unit of charge. Otherwise, we charge s as follows:

- C1. s only intersects one rectangle $s' \in S'$. Charge s to s' . By condition (A3), each rectangle in S' receives at most one unit of charge.
- C2. s intersects s' in a corner (i.e., s contains a vertex of s'). Charge s to s' . Since I and S' are independent sets, each rectangle in S' receives at most four units of charge.

Let $I' \subseteq I$ be the set of rectangles of I that have *not* yet been charged. Thus far each rectangle in S' has received at most 5 charges, so therefore $|I \setminus I'| \leq 5|S'|$. We now bound the number of rectangles in I' . I' has the following property: each rectangle in I' intersects exactly two rectangles in S' (see Figure 5.1(b)); C1 and C2 above deal with all other intersection possibilities. We map each rectangle in S' and I' to a vertex and an edge, respectively, by constructing a bipartite multi-graph $G_{I'} = (V, E)$ where each vertex in V corresponds to a rectangle in S' , and there is an edge between v_j and v_k if there is a rectangle $s \in I'$ that intersects both s_{i_j} and s_{i_k} . Clearly each rectangle in I' maps to an edge in $G_{I'}$.

We now show that $G_{I'}$ is a planar multi-graph as follows. Replace each rectangle $s' \in S'$ with its center $c(s')$, and each rectangle $s \in I'$ with a polygonal curve (consisting of three piecewise-linear segments) as illustrated in Figure 5.1(c). It is clear, given the independence property of the rectangles in I' and S' , that the above embedding is non-intersecting. Hence $G_{I'}$ is a planar multi-graph.

For a planar graph, the number of edges is at most thrice the number of vertices. Since $G_{I'}$ is a planar multi-graph with at most two edges between any pair of vertices by condition (A4), we have $|I'| = |E| \leq 6|V| = 6|S'|$. Hence, combining the bounds,

$$|I| = |I'| + |I \setminus I'| \leq 6|S'| + 5|S'| = 11|S'|.$$

□

Lemma 5.4. *Given a set S of n rectangles in \mathbb{R}^2 such that no two rectangles pierce, let G_{np} denote the corresponding (non-piercing) intersection graph. Then $\omega(G_{np}) \cdot \alpha(G_{np}) \geq \frac{n}{16}$, and a c_0 -approximation to the maximum independent set can be computed in polynomial time,*

Proof. First, by Turán's Theorem [PA95], there exists an independent set of size at least $n^2/4|E|$. Second, for each edge $e = \{s_i, s_j\}$, due to the properties of non-piercing graphs, either s_i contains

a vertex of s_j or vice versa. Define c_i (a lower-bound on the number of times rectangle s_i 's vertex is contained in another rectangle) as follows. If s_i contains a vertex of s_j , charge the edge e to vertex s_j by incrementing c_j , otherwise increment c_i . Clearly, $\sum_{i=1}^n c_i = |E|$, and hence there exists a rectangle s_k whose vertices are contained in more than $|E|/n$ rectangles, and therefore at least one vertex contained in more than $|E|/4n$ rectangles. These rectangles share a common point, and thus form a clique of size greater than $|E|/4n$. Hence,

$$\omega(G_{np}) \cdot \alpha(G_{np}) \geq \frac{n^2}{4|E|} \cdot \frac{|E|}{4n} = \frac{n}{16}.$$

We describe an iterative algorithm to compute a r -maximal set, which, by Lemma 5.3, is a constant-factor approximation of the maximum independent set of G_S . In the beginning of the i -th iteration the algorithm maintains a set S_i of i pairwise-disjoint rectangles. In the i -th iteration, the algorithm checks whether conditions (A2)–(A4) are satisfied. If the answer is yes, we return S_i , as it is a r -maximal set. Otherwise,

1. Condition (A2) violated. Then there exists a set $s \in S \setminus S'$ that does not intersect any $s' \in S'$. Set $S_{i+1} = S_i \cup \{s\}$.
2. Condition (A3) violated. Then there exist two disjoint rectangles $s, t \in S \setminus S'$ that intersect some $s' \in S'$ but no other rectangle in S' . Set $S_{i+1} = S_i \setminus \{s'\} \cup \{s, t\}$.
3. Condition (A4) violated. Then there exist three pairwise-disjoint rectangles $s, t, u \in S \setminus S'$ that all intersect $s', t' \in S'$, but no other rectangle in S' . Set $S_{i+1} = S_i \setminus \{s', t'\} \cup \{s, t, u\}$.

It is clear that S_i is a set of pairwise-disjoint segments. The process can continue for at most n iterations, and the final set S_j is obviously a r -maximal set. Since each of the conditions (A2)–(A4) can be checked in polynomial time, the total running time is polynomial. \square

General Rectangle Intersection Graphs

Combining Lemma 5.2 and Lemma 5.4, we attain the following.

Lemma 5.5. *Given a set S of n rectangles in \mathbb{R}^2 , an independent set of size at least $\frac{\alpha(S)}{d_0 \cdot \omega(S)}$, for some constant $d_0 \leq 16c_0$, can be computed in polynomial time.*

Proof. By Lemma 5.2, compute the maximum independent set, say $A \subseteq S$, in the piercing graph of S . Note the following:

- (i) $|A| \geq \alpha(S)$ as an independent set in S is an independent set in piercing graph of S , and
- (ii) $\omega(A) \leq \omega(S)$, since a clique in A is a clique in S .

By Lemma 5.4, $\alpha(A) \cdot \omega(A) \geq |A|/16$. Hence, using (i) and (ii) above, it follows that $\alpha(A) \geq |A|/16\omega(A) \geq \alpha(S)/16\omega(S)$. Since the intersection graph of A is non-piercing, Lemma 5.4 gives an algorithm which returns an independent set of size at least $\alpha(A)/c_0 \geq \alpha(S)/16c_0\omega(S)$. \square

Remark: Independently of our work, Lewin-Eytan *et al.* [LENO02] also proved Lemma 5.5 using considerably more complicated LP-rounding techniques.

From now on, let $\alpha(S) = \beta n$, for some $\beta \leq 1$.

Theorem 5.6. *Given S with $\alpha(S) = \beta n$, one can compute an independent set of size $\frac{\alpha(S)}{4d_0 \cdot (1/\beta)}$ in polynomial time.*

Proof. The algorithm repeatedly extracts large cliques (one can compute the maximum clique in rectangle intersection graphs in polynomial time [IA83]) until a good independent set is found, as follows. Set $S_0 = S$, and let S_i be the set of rectangles in the i -th iteration. For now assume that the value of β is known. At the i -th iteration, if there exists a clique \mathcal{C} of size at least $2/\beta$, remove \mathcal{C} from S_i , i.e., set $S_{i+1} = S_i \setminus \mathcal{C}$, and reiterate. If no such clique exists, compute the maximum independent set, say A , in the piercing graph of S_i , and return the maximum independent set in A . Assume the algorithm stops after j iterations, i.e. $\omega(S_j) \leq 2/\beta$. Note that $j \leq n/(2/\beta) = \alpha(S)/2$. Since at most one rectangle from the independent set can be in a clique, each iteration removes at most one rectangle from the optimal independent set of S , hence $\alpha(S_j) \geq \alpha(S) - j$. From Lemma 5.5, one can thus compute an independent set of size at least

$$\frac{\alpha(S_j)}{d_0 \omega(S_j)} \geq \frac{\alpha(S) - j}{d_0(2/\beta)} \geq \frac{\alpha(S) - \alpha(S)/2}{d_0(2/\beta)} = \frac{\alpha(S)}{(4d_0/\beta)},$$

yielding the desired result. Note that we do not know the value of β , but can run the above algorithm for all the n possible values, and return the maximum. \square

This page was intentionally left blank

6. A PTAS for the Geometric Hitting-Set Problem for Disks

Different mathematicians study papers in different ways, but when I read a mathematical paper in a field in which I'm conversant, I concentrate on the thoughts that are between the lines. I might look over several paragraphs or strings of equations and think to myself, "Oh yeah, they're putting in enough rigamarole to carry such-and-such an idea." When the idea is clear, the formal setup is usually unnecessary and redundant. I often feel that I could write it out myself more easily than figuring out what the authors actually wrote. It's like a new toaster that comes with a 16-page manual. If you already understand toasters and if the new toaster looks like previous toasters you've encountered, you might just plug it in and see if it works, rather than first reading all the details in the manual.

William Thurston

In this chapter we will consider the following hitting-set problem: given a set P of n points, and set systems where the ranges in \mathcal{D} are induced by various geometric objects, compute the minimum-sized hitting-set for \mathcal{D} . Specifically, we show that:

- Given a set P of n points, and a set \mathcal{H} of m half-spaces in \mathbb{R}^3 , one can compute a $(1 + \varepsilon)$ -approximation to the smallest subset of P that hits all the half-spaces in \mathcal{H} in $O(mn^{O(\varepsilon^{-2})})$ time.
- Given a set P of n points in \mathbb{R}^2 , and a set of r -admissible regions \mathcal{D} , one can compute a $(1 + \varepsilon)$ -approximation to the smallest subset of P that hits all the regions in \mathcal{D} in $O(mn^{O(\varepsilon^{-2})})$ time. This includes pseudo-disks (they are 2-admissible), same-height rectangles, circular disks, translates of convex objects etc. See Definition 6.1.1 for the definition of an r -admissible set of regions.

Our algorithms are based on the *local-search* paradigm, the basic idea of which is the following: start with any feasible solution, and repeatedly improve the current solution by iterating over local improvement steps. While local-search has been used quite effectively as a practical heuristic, there are very few cases where one can get provable guarantees on the resulting solution. These are even rarer for algorithms for geometric problems.

Our algorithm for both the problems is the following simple local search algorithm: start with any hitting set $S \subseteq P$ (e.g., take all the points of P), and iterate local-improvement steps of the

following kind: If any k points of S can be replaced by $k - 1$ points of P such that the resulting set is still a hitting set, then perform the swap to get a smaller hitting set. Halt if no such local improvement is possible. We will call this a k -level local search algorithm. Then our main result is the following:

Theorem 6.1. *Let P be a set of n points in \mathbb{R}^3 (resp. \mathbb{R}^2), and let \mathcal{H} (resp. \mathcal{D}) be the geometric objects as above. There is a universal constant c such that for any $\epsilon > 0$, a $(c/\epsilon)^2$ -level local search algorithm returns a hitting set of size at most $(1 + \epsilon) \cdot \text{OPT}$, where OPT is the size of an optimal (smallest) hitting set.*

An important corollary of the above result for halfspaces in \mathbb{R}^3 is the following:

Corollary 6.0.1. *Given a set P of points and a set \mathcal{D} of disks in the plane, there exists a PTAS for approximating the smallest subset of the disks which cover all points in P .*

Proof. We map each point (p, q) in the plane to the halfspace $z - 2px - 2qy + p^2 + q^2 \leq 0$ and each disk $(x - a)^2 + (y - b)^2 \leq r^2$ to the point $(a, b, a^2 + b^2 - r^2)$. It can be checked that this map preserves incidence relations between points and disks. Now, the PTAS for hitting sets for halfspaces in \mathbb{R}^3 gives a PTAS for our problem. \square

Note that, for any fixed k , the naive implementation of the k -level local search algorithm takes polynomial time: start the algorithm with the entire set P as (the most likely sub-optimal) hitting set P' . The size of P' decreases by at least one at each local-improvement step. Hence, there can be at most n steps of local improvement, where there are at most $\binom{n}{k} \cdot \binom{n}{k-1} \leq n^{2k-1}$ different local improvements to verify. Checking whether a certain local improvement is possible takes $O(nm)$ time. Hence the overall running time of the algorithm is $O(mn^{2k+1})$.

Combinatorial bounds on ϵ -nets via Local Search. As a side result, we show that the local search technique can also be used to prove the existence of small-sized ϵ -nets. Specifically, we show that for the case where we have points in the plane and ranges consist unit squares in the plane, a simple local-search method gives the optimal bound of $O(1/\epsilon)$ for the size of the ϵ -net. It is quite easy to prove the same result using other techniques but it is interesting that the local search technique can be used to prove this. This kind of result is currently known only for half-spaces in \mathbb{R}^2 and is implied by the proof of the existence of $O(1/\epsilon)$ size ϵ -nets by Pach and Woeginger [PW90]. It is not at all clear that the same holds for half-spaces in \mathbb{R}^3 . We conjecture that this holds for more general range spaces defined by a set of points and an r -admissible set of regions in the plane – we leave this as an open problem.

6.1 A $(1 + \epsilon)$ -approximation scheme for hitting geometric sets

Let $\mathcal{R} = (P, \mathcal{D})$ be a range space where P is the *ground set* and $\mathcal{D} \subseteq 2^P$ is the set of ranges. A minimum hitting set for \mathcal{R} is a subset $Q \subseteq P$ of the smallest size such that $Q \cap D \neq \emptyset$, for all $D \in \mathcal{D}$. In this section we will show that given any parameter $\epsilon > 0$, a $O(\epsilon^{-2})$ -level local search returns a hitting set whose size is at most $(1 + \epsilon)$ times the size of the minimum hitting set for range spaces that satisfy the following *locality condition*.

Locality Condition. A range space $\mathcal{R} = (P, \mathcal{D})$ satisfies the locality condition if for any two disjoint subsets $R, B \subseteq P$, it is possible to construct a planar bipartite graph $G = (R, B, E)$ with all edges going between R and B such that for any $D \in \mathcal{D}$ with $D \cap R \neq \emptyset$ and $D \cap B \neq \emptyset$, there exist two vertices $u \in D \cap R$ and $v \in D \cap B$ such that $(u, v) \in E$.

For example, if P is a set of points in the plane and \mathcal{D} is defined by intersecting P with a set of circular disks, then $\mathcal{R} = (P, \mathcal{D})$ satisfies the locality condition. To see this consider, for any given R and B , the Delaunay triangulation G of $R \cup B$. Removing the non red-blue edges from the triangulation gives the required bipartite planar graph since for each disk D in the plane, the vertices in $(R \cup B) \cap D$ induce a connected subgraph of G and hence there must be an edge between a vertex in $D \cap R$ and a vertex in $D \cap B$ whenever both the intersections are non-empty.

The fact that the delaunay graph G of a set of points P has the property that the subgraph of G induced by the subset of P lying in an arbitrary disk is connected, can be easily proved by contradiction. If there are disks for which this is not the case then consider a disk D of smallest radius which contains two vertices u and v belonging to different connected components of the subgraph induced by $D \cap P$. It is not hard to see that both u and v must lie on the boundary of D for otherwise we can shrink D to find a disk $D' \subset D$ which still contains u and v . Now, if D does not have any point of P in its interior then (u, v) is a delaunay edge which contradicts the assumption that u and v belong to different components. On the other hand, if there is a point w inside D then either u and w belong to different components or v and w belong to different components. In either case, we can shrink D to get $D' \subset D$ which still has two vertices belonging to different connected components of $D' \cap P$, thus contradicting the minimality of D .

Let us now return to the hitting set problem. For any vertex v in a graph G , denote by $N_G(v)$ the set of neighbors of v . Similarly, for any subset of the vertices W of G , let $N_G(W)$ denote the set of all neighbors of the vertices in W , i.e., $N_G(W) = \bigcup_{v \in W} N_G(v)$. Our basic theorem is the following:

Theorem 6.2. *Let $\mathcal{R} = (P, \mathcal{D})$ be a range space satisfying the locality condition. Let $R \subseteq P$ be an optimal hitting set for \mathcal{D} , and $B \subseteq P$ be the hitting set returned by a k -level local search. Furthermore, assume $R \cap B = \emptyset$. Then there exists a planar bipartite graph $G = (R, B, E)$ such that for every subset $B' \subseteq B$ of size at most k , $|N_G(B')| \geq |B'|$.*

Proof. Let $\mathcal{R} = (P, \mathcal{D})$ be a range space satisfying the locality condition where P is set of size n and \mathcal{D} is a set of m subsets of P . From now on, we will call R and B the red points and the blue points respectively. Since no local improvement is possible in B , we can conclude that no k blue points can be replaced by $k - 1$ or fewer non-blue points. In particular, no k blue points can be replaced by $k - 1$ or fewer red points.

Let G be the bipartite planar graph between R and B , given by the locality condition for \mathcal{R} . Since both R and B are hitting sets for \mathcal{R} , we know that each range in \mathcal{D} has both red and blue points.

Claim 6.3. *For any $B' \subseteq B$, $(B \setminus B') \cup N_G(B')$ is a hitting set for \mathcal{R} .*

Proof. If there is range $D \in \mathcal{D}$ which is only hit by the blue points in B' , then one of those blue points has a red neighbor that hits D and therefore $N_G(B')$ hits D . Otherwise, D is hit by some point in $B \setminus B'$. \square

This finishes the proof, since the above claim implies that if $B' \subseteq B$ is a set of at-most k blue points, then $|N_G(B')| \geq |B'|$ since otherwise a local improvement would be possible in B . \square

Note that we can always assume, without loss of generality, that $B \cap R = \emptyset$. If not, let $I = B \cap R$, $P' = P \setminus I$, $B' = B \setminus I$, $R' = R \setminus I$ and let \mathcal{D}' be the set of ranges that are not hit by the points in I . B' and R' are disjoint. Also, R' is a hitting set of minimum size for the hitting set problem with points P' and the ranges in \mathcal{D}' . If we can show that $|B'|$ is approximately equal to $|R'|$, we can conclude that $|B|$ is approximately equal to $|R|$.

Now, the following lemma (also proved independently in Har-Peled and Chan [CHP09]) implies that given any parameter ϵ , a k -level local search with $k = c^2\epsilon^{-2}$ gives a $(1 + \epsilon)$ -approximation to the minimum hitting set problem for \mathcal{R} .

Lemma 6.4. *Let $G = (R, B, E)$ be a bipartite planar graph on red and blue vertex sets R and B , $|R| \geq 2$, such that for every subset $B' \subseteq B$ of size at most k , where k is a large enough number, $|N_G(B')| \geq |B'|$. Then $|B| \leq (1 + c/\sqrt{k}) |R|$, where c is a constant.*

The above lemma follows directly from Lemma 1 of Frederickson [Fre87], which is a refinement of the Lipton-Tarjan separator theorem [LT77]. We state it in a slightly different way below.

Theorem 6.5 (Planar graph partition with small boundary size [Fre87]). *Given a planar graph H with n vertices and a parameter t , the vertices of H can be divided into groups of size at most t so that, for each edge there is a group containing both its end points and the total number of vertices of a group shared with other groups, summed over all groups, is at most $\gamma n/\sqrt{t}$, where γ is a fixed constant.*

Note that some vertices belong to more than one group – these vertices are called *boundary* vertices. Furthermore, each non-boundary vertex has edges only to members of its own group (which could include some boundary vertices).

Proof of Lemma 6.4. Let $r = |R|$ and $b = |B|$. Consider the groups of G formed according to Theorem 6.5 with the parameter $t = k$. Each group has at most k vertices. Consider the i^{th} group and let r_i^∂ and b_i^∂ be the number of red and blue boundary vertices respectively in the group. Similarly, let b_i^{int} and r_i^{int} be the number of red and blue interior (non-boundary) vertices in this group. Theorem 6.5 guarantees that $\sum_i r_i^\partial + b_i^\partial \leq \gamma(r + b)/\sqrt{k}$. Since there are at most k interior blue vertices in the group, by the expansion condition of the theorem, their neighborhood must be at least as large as their own number, i.e., $b_i^{\text{int}} \leq r_i^{\text{int}} + r_i^\partial$. Adding b_i^∂ to both sides and summing over all i we have

$$\begin{aligned} b &\leq \sum_i (b_i^{\text{int}} + b_i^\partial) \leq \sum_i r_i^{\text{int}} + \sum_i (r_i^\partial + b_i^\partial) \\ &\leq r + \gamma(r + b)/\sqrt{k} \end{aligned}$$

Let us assume that $k \geq 4\gamma^2$ and set $c = 4\gamma$. Then,

$$\begin{aligned}
b &\leq r \frac{1 + \gamma/\sqrt{k}}{1 - \gamma/\sqrt{k}} = r(1 + \gamma/\sqrt{k})(1 + (\gamma/\sqrt{k}) + (\gamma/\sqrt{k})^2 + \dots) \\
&\leq r(1 + \gamma/\sqrt{k})(1 + 2\gamma/\sqrt{k}) \quad (\text{since } \gamma/\sqrt{k} \leq 1/2) \\
&= r(1 + 3\gamma/\sqrt{k} + 2\gamma^2/k) \\
&\leq r(1 + 4\gamma/\sqrt{k}) \quad (\text{since } 2\gamma^2/k \leq \gamma/\sqrt{k}) \\
&= r(1 + c/\sqrt{k}).
\end{aligned}$$

□

PTAS for an r -admissible set of regions. It turns out that the locality condition, by a more complicated construction of the planar graph G [PR08], also holds for an r -admissible set of regions, for any r , in the plane. This yields a PTAS for the minimum hitting set problem with an r -admissible set of regions in the plane. The definition of an r -admissible set of regions is as follows:

Definition 6.1.1. *A set of regions in \mathbb{R}^2 , each of which is bounded by a closed Jordan curve, is called r -admissible (for r even), if for any two s_1, s_2 of the regions, the Jordan curves bounding them cross in $l \leq r$ points, (for some even l), and both $s_1 \setminus s_2$ and $s_2 \setminus s_1$ are connected regions.*

As mentioned earlier, this includes pseudo-disks (they are 2-admissible), same-height rectangles, circular disks, translates of convex objects etc.

PTAS for half-spaces in \mathbb{R}^3 . Given a set of half-spaces and a set of points in \mathbb{R}^3 , we first pick one of the points o and add it to our hitting set. We then ignore o and all half-spaces containing it. Let $\mathcal{R} = (P, \mathcal{D})$ be the range space defined by the remaining set of points and the remaining set of half-spaces. A PTAS for \mathcal{R} gives a PTAS for the original problem. We will show that \mathcal{R} satisfies the locality condition. Let R and B be disjoint red and blue subsets of P .

We construct the required graph G on the vertices $R \cup B$ in two stages and prove its planarity by giving its embedding on the boundary $\partial\mathcal{C}$ of the convex hull \mathcal{C} of $R \cup B$. In the first stage, we add all red-blue edges (1-faces) of \mathcal{C} to G . In the second stage we map each red or blue point p lying in the interior \mathcal{C} to a triangular face $\Delta(p)$ of \mathcal{C} that intersects the ray op emanating from o and passing through p .¹

Let Q be the set of points mapped to a triangle Δ . We will construct a planar bipartite graph on Q and the corners of Δ and embed it so that the edges lie inside Δ . If Δ has two red corners and one blue corner, we add an edge between each red point in Q to the blue corner of Δ and each blue point of Q to the two red corners of Δ . It is quite easy to see that this can be done so that the graph remains planar. The case when Δ has two blue corners and one red corner is handled similarly. Consider now the case when all corners of Δ are red and let r_1, r_2 and r_3 be the corners. In this case we will connect at most one blue point of Q to all three corners of Δ and we will connect the rest of the blue points to two of the corners of Δ . Again, it is clear that this can be done regardless of the selected corners while keeping the graph planar.

¹Here we are assuming that each face of \mathcal{C} is a triangle, since one can always triangulate the faces.

For each blue point $b \in Q$, we try to find a corner c of Δ such that there is no half-space $h \in \mathbb{R}^3$ that contains only b and c out of $(B \cap Q) \cup \{r_1, r_2, r_3, o\}$. If we can find such a corner c , then we put an edge between b and the two corners of Δ other than c . There can be at most one blue point in Q for which we cannot find such a corner and we will connect that blue point to all three corners of Δ . For contradiction, assume that there are two points b_1 and b_2 in Q such that for each pair of red and blue points in $F = \{r_1, r_2, r_3, b_1, b_2\}$ there is a half-space in \mathbb{R}^3 containing exactly those two points of F . This means that each $r_i b_j$ is an edge in the convex hull of F and therefore F is in convex position. The Radon partition [Mat02] of F is then a $(3, 2)$ -partition. Since the blue points lie on the same side of the plane containing Δ , the partition with two points has one red point and one blue point and there cannot be a half-space containing exactly these two among the points of F , contradicting our assumption. The case when Δ has three blue corners is handled similarly. The construction of G is complete.

We now show that for any half-space $h \in \mathbb{R}^3$ that does not contain o and contains both red and blue points, there is an edge in G between a red point and blue point both of which lie in h . If h contains both red and blue points which lie on $\partial\mathcal{C}$ then there is a red-blue edge among two of those due to the edges added in the first stage. Otherwise assume, without loss of generality, that only the red points in h lie on $\partial\mathcal{C}$. Consider the half-space h' parallel to and contained in h which contains the smallest number of points and still contains both red and blue points. h' contains exactly one blue point b . Since h , and hence h' , does not contain o , h' must contain one of the corners of the triangle Δ that b is mapped to. If b is connected to all three corners of Δ in G , we are trivially done. Also, if h contains two of the corners of Δ , then we are done since b is connected to at least one of those corners. If h' contains exactly one corner c of Δ , then b must be adjacent to c by the way we constructed the graph G . Hence, in all cases, b is connected to one of the red points in h' .

6.2 Combinatorial Bounds on ε -nets via Local Search

Consider the range space $\mathcal{R} = (P, \mathcal{D})$ in which P is a set of points in the plane and \mathcal{D} is defined by intersecting P with a set of unit squares in the plane. Construct an ε -net for \mathcal{R} , say Y , using the 3-level local search: starting with $Y = P$, keep improving Y as long as there exists a subset of size at most three of Y that can be swapped to get a smaller set. We now argue that $|Y| = O(1/\varepsilon)$.

For the argument we will consider an equivalent problem. We will replace each of the squares by a point at its center and each of the points with a unit square centered at it. The task now is to pick the smallest subset of the squares which cover all points which are covered by more than an ε fraction of the squares. Let the number of squares be n and the number of points be m . We will refer to the set of squares corresponding to points in P by S and the set of squares corresponding to the points in Y by M .

First some definitions. Call the squares in M the “ ε -net squares” and the squares in $S \setminus M$ as “normal squares”. A point $p \in \mathbb{R}^2$ is *dense* if it is covered by more than εn squares in S . Each $s \in M$ must have a *personal dense point*, i.e., a dense point which no other square in M covers. Fix any unit gridding of the plane, and call a grid point p *active* if at least one of the four cells touching it contains a dense point. Denote the set of active grid points by A . The following claim is easy to show.

Claim 6.6. $|A| = O(1/\varepsilon)$.

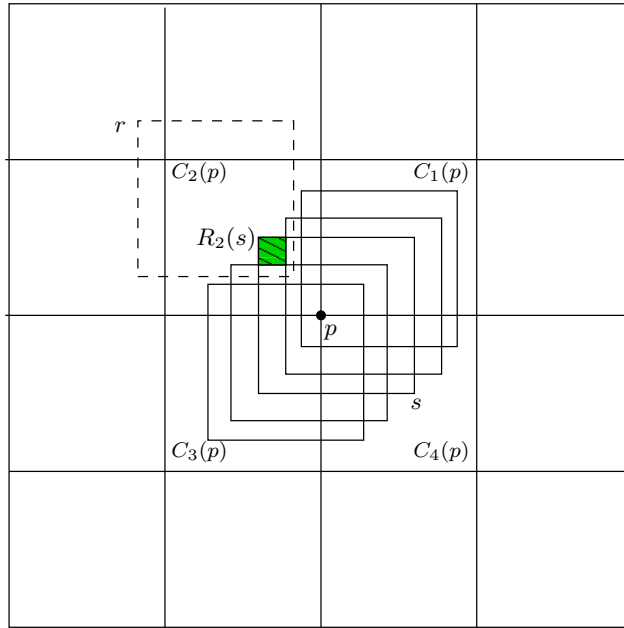


Figure 6.1: The normal square r covers the ϵ -net square s and stabs its neighbors (in the cascade $M_2(p)$) in the cell $C_2(p)$.

Proof. By a packing argument, each active point has ϵn unit squares intersecting one of its four adjacent squares. These squares contribute a constant number of active points, and there can be only $O(1/\epsilon)$ such sets. \square

Each unit square $s \in S$ contains exactly one of the grid points, and for the squares in M , this grid point belongs to A . For each active grid point $p \in A$, label the four cells around it as $C_1(p)$, $C_2(p)$, $C_3(p)$ and $C_4(p)$ in counter-clockwise order. For each cell $C_i(p)$, refer to its opposite cell as $C'_i(p)$ (e.g., $C_1(p)$ is the opposite cell to $C_3(p)$). Denote the set of squares in M that contain the grid point p by $M(p)$, and among these, those that have a personal dense point in $C_i(p)$ as $M_i(p)$. Each square of M containing p must belong to at least one of the four $M_i(p)$'s. Each set $M_i(p)$ forms a *cascade* and there is a natural linear order on them. Call the squares which are not the first or the last in this order the *middle* squares of $M_i(p)$. Each square $s \in M_i(p)$ has some region in $C_i(p)$ which is not covered by other squares in $M_i(p)$ and we denote this region by $R_i(s)$ (see Figure 6.1). This square s also has a region in $C'_i(p)$ which is not covered by other squares in $M_i(p)$, denoted by $R'_i(s)$. For a normal square r and an ϵ -net square $s \in M_i(p)$ we say that “ r stabs s in $C_i(p)$ ” if r intersects the region $R_i(s)$ and we say that “ r covers s in $C_i(p)$ ” if r contains the region $R_i(s)$. Note that if r covers s then r also stabs s .

Lemma 6.7. *No three middle squares in $M_i(p)$ have a common coverer in both $C_i(p)$ and $C'_i(p)$. Furthermore, no five squares in $M_i(p)$ are stabbed by a common square in both $C_i(p)$ and $C'_i(p)$. Also, no four squares in $M_i(p)$ which have a personal dense point only in $C_i(p)$ are stabbed by a common square in $C_i(p)$.*

Proof. If three middle squares in $M_i(p)$ have a common coverer r in $C_i(p)$ and a common coverer

r' in $C'_i(p)$, then a local improvement is possible by replacing the three squares by two squares r and r' in the ϵ -net. Similarly, if five squares are stabbed by a common square r (resp. r') in $C_i(p)$ (resp. $C'_i(p)$), then the three middle squares among them are covered by r (resp. r'), which is not possible by the first statement. If four squares in $M_i(p)$ which have a personal dense point only in $C_i(p)$ are stabbed by a square r then the two middle squares among the four can be replaced by r contradicting the assumption that no local improvement is possible. \square

For any square $s \in M$, let $N(s)$ be the set of normal squares intersecting s . Also, let $Z(p) = \cup_{s \in M(p)} N(s)$ be the neighborhood of $M(p)$ and $Z_i(p) = \cup_{s \in M_i(p)} N(s)$ be the neighborhood of $M_i(p)$.

Claim 6.8. $|M(p)| \leq 28 \frac{|Z(p)|}{\epsilon n} + 27$.

Proof. Since $M(p) = \cup_i M_i(p)$, we have that either $|M_1(p) \cup M_3(p)| \geq |M(p)|/2$ or $|M_2(p) \cup M_4(p)| \geq |M(p)|/2$. Let us assume without loss of generality that $|M_1(p) \cup M_3(p)| \geq |M(p)|/2 \geq 14 \lfloor |M(p)|/28 \rfloor$. Set $M' = M_1(p) \cap M_3(p)$ and let $t = |M'|$. Now, there are two cases depending on whether t is more than $8 \lfloor |M(p)|/28 \rfloor$. Suppose first that it is not. Then there are at least $6 \lfloor |M(p)|/28 \rfloor$ squares which have a personal dense point in either C_1 or C_3 but not both. Assume, without loss of generality, that half of them have a personal dense point in C_1 but not C_3 . Then, since each of these personal dense points are covered by at least ϵn squares and, by the last statement of Lemma 6.7, no square can cover more than 3 personal dense points, we have that $Z(p) \geq \lfloor |M(p)|/28 \rfloor \cdot \epsilon n$.

Now suppose that t is at least $8 \lfloor |M(p)|/28 \rfloor$. The squares in M' have a personal dense point in both $C_1(p)$ and $C_3(p)$. and let s_1, s_2, \dots, s_t be the squares of M' along the cascade defined by them. For each square s_j , define its *red (blue) successor* to be the square s_k with the smallest index $k > j$ such that s_j and s_k are not stabbed by a common square in $C_1(p)$ ($C_3(p)$). Note that a square may not have a red or blue successor. Let us also say that a red or blue successor of a square s_i is *far* if the successor is s_j with $j - i \geq 5$ and *near* otherwise. If some square s_i has a red (blue) successor s_j that is far then s_i the squares of M' between s_i and s_{j-1} , of which there are at least 5, are stabbed by a common square in $C_1(p)$ ($C_3(p)$). Lemma 6.7 therefore implies that both red and blue successors of a square cannot be far. At least one of them has to be near. Assume, without loss of generality, that at least half of the squares in M' have a red successor that is near. Let M'' be the set of such squares. Let M''' be the set of squares in which we take every fourth square of M'' starting with the first in the cascade defined by them. Clearly no two squares in M''' are stabbed by a common square in $C_1(p)$ since otherwise one of them would have a far red successor. Now, since $|M'''| \geq \lfloor |M(p)|/28 \rfloor$ and each normal square can contain the personal dense point of at most one of the squares of M''' in $C_1(p)$, we again have that $|Z(p)| \geq \lfloor |M(p)|/28 \rfloor \cdot \epsilon n$. The claim follows. \square

We now bound the size of the ϵ -net M .

Claim 6.9. $|M| = O(1/\epsilon)$.

Proof. A square can belong to the neighborhood of at most nine active points, i.e., $\sum_{p \in A} |Z(p)| \leq 9n$. Summing the inequality in Claim 6.8 over all $p \in A$ and using Claim 6.6, one gets the required

statement: $|M| = \sum_{p \in A} |M(p)| = \frac{\sum_{p \in A} |Z(p)|}{\epsilon n} + 27|A| = O(1/\epsilon)$.

□

This page was intentionally left blank

Part III

Data Depth

7. A Generalization of the Centerpoint Theorem

Thought is only a flash in the middle of a long night, but the flash that means everything.

Henri Poincaré

Given a set $P = \{p_1, \dots, p_n\}$ of n points in \mathbb{R}^d and a finite set $Q \subset \mathbb{R}^d$, define the following:

$$\varepsilon(P, Q) = \min\{\varepsilon \mid |\mathcal{C} \cap Q| \neq \emptyset \ \forall \text{ convex sets } \mathcal{C} \text{ s.t. } |\mathcal{C} \cap P| > \varepsilon n\}$$

and let $\varepsilon_i^d(P) = \min_{Q, |Q|=i} \varepsilon(P, Q)$. Set $\varepsilon_i^d = \sup_P \varepsilon_i^d(P)$. In other words, given any set P of n points in \mathbb{R}^d , the set of all convex sets containing $\varepsilon_i^d n$ points of P can be hit by i points. These i points are said to form a *weak* ε_i^d -net for P . Recall that the centerpoint theorem in d dimensions states that $\varepsilon_1^d = \frac{d}{d+1}$.

In this chapter we prove the following result:

Theorem 7.1. *Given a set P of n points in \mathbb{R}^d , and integers $r, s \geq 0$,*

$$\varepsilon_{r+ds+1}^d \leq \frac{\varepsilon_r^d \cdot (1 + (d-1)\varepsilon_s^d)}{1 + \varepsilon_r^d \cdot (1 + (d-1)\varepsilon_s^d)}$$

where we define $\varepsilon_0^d = 1$.

Fix a direction $\vec{u} \in \mathbb{R}^d$ which we call the *upward* direction. For a point $p \in \mathbb{R}^d$, let $f_u(p) = \langle u, p \rangle$ denote the *height* of the point p in the upward direction ($\langle u, p \rangle$ denotes the inner product of u and p). For a convex set \mathcal{C} , let $f_u(\mathcal{C})$ denote the height of the lowest point in \mathcal{C} , i.e. $f_u(\mathcal{C}) = \inf_{p \in \mathcal{C}} f_u(p)$.

Construction. Let $a, b \in [0, 1]$ be two reals to be fixed later.

Let $\mathcal{H} = \{h_1, \dots, h_k\}$ be the set of all closed halfspaces which contain at least an points of P and whose bounding hyperplane passes through d points in P . Define $\mathcal{H}^d = \{(h_{i_1}, h_{i_2}, \dots, h_{i_d}) \mid |P \cap (h_{i_1} \cap h_{i_2} \dots \cap h_{i_d})| \geq bn, \text{ where } h_{i_1}, h_{i_2}, \dots, h_{i_d} \in \mathcal{H}\}$ to be the set of all d -tuples of halfspaces in \mathcal{H} whose intersection contains at least bn points of P . Consider the d -tuple, say $(h_{l_1}, \dots, h_{l_d})$, such that

1. $(h_{l_1}, \dots, h_{l_d}) \in \mathcal{H}^d$




Figure 7.1: Illustration of Theorem 7.1

2. $(h_{l_1} \cap \dots \cap h_{l_d})$ has the highest lowest-intersection point among the d -tuples of halfspaces in \mathcal{H}^d , i.e., $f_u(h_{l_1} \cap \dots \cap h_{l_d}) = \max_{(h_{i_1}, \dots, h_{i_d}) \in \mathcal{H}^d} f_u(h_{i_1} \cap \dots \cap h_{i_d})$

We choose the upward direction \vec{u} so that the d -tuple $(h_{l_1}, \dots, h_{l_d})$ is well defined. Note that $f_u(h_{i_1} \cap \dots \cap h_{i_d}) = -\infty$ iff $h_{i_1} \cap \dots \cap h_{i_d}$ is unbounded in the downward direction $-\vec{u}$. Let \mathcal{P} be the convex hull of P and let h_{j_1}, \dots, h_{j_d} be d halfspaces defining a vertex v of \mathcal{P} and containing P . Choose the upward direction \vec{u} so that the vertex v is the unique lowest vertex of the polyhedron $\mathcal{P}' = h_{j_1} \cap \dots \cap h_{j_d}$ in the upward direction and each of the points $p \in P$ get a unique height. Such a choice of \vec{u} ensures that the bounding hyperplane of no halfspace in \mathcal{H} has a normal parallel to the upward direction u and there is at least one d -tuple of halfspaces in \mathcal{H}^d whose intersection is bounded in the downward direction $-\vec{u}$. Therefore, $(h_{l_1}, \dots, h_{l_d})$ is well defined and the lowest point in $h_{l_1} \cap \dots \cap h_{l_d}$ is unique.

Let \mathcal{R} be the polyhedron $\{h_{l_1} \cap \dots \cap h_{l_d}\}$. Without loss of generality, we can assume that \mathcal{P} is full-dimensional and hence \mathcal{R} is full-dimensional. Let \mathcal{R}_{l_i} be the intersection of the halfspaces in $\{h_{l_1}, \dots, h_{l_d}\}$ except l_i i.e., $\mathcal{R}_{l_i} = \bigcap_{k \in [1, d], k \neq i} h_{l_k}$. Since each of the halfspaces contain at least an points from P , $|P \cap \mathcal{R}_{l_i}| \geq (d-1)an - (d-2)n$. Construct and return the set $Q = \{x\} \cup Q' \cup Q_{l_1} \cup \dots \cup Q_{l_d}$, where

1. x is the unique lowest point in $h_{l_1} \cap \dots \cap h_{l_d}$.
2. Q' is an ε_r^d -net for the point set $P \setminus (P \cap h_{l_1} \cap \dots \cap h_{l_d})$ using r points.
3. Q_{l_i} is an ε_s^d -net for the point set $P \setminus (P \cap \mathcal{R}_{l_i})$ using s points.

Lemma 7.2. Q is an a -net for P , and has size $r + ds + 1$.

Proof. The size of Q is obvious from the construction, and we show that it is an a -net for the value required in the statement of the theorem. We first need the following crucial fact.

Claim 7.3. Let C' be a convex set containing at least an points of P which does not contain x and contains points from $P \cap h_{l_1} \cap \dots \cap h_{l_d}$. Then, $|P \cap C' \cap \mathcal{R}_{l_i}| < bn$ for some $i \in [1, d]$.

Proof. For contradiction, assume that C' intersects all R_{l_i} at at least bn points of P . Let \mathcal{R}' be the convex hull of $P \cap C'$. Then, \mathcal{R}' does not contain x , and therefore there exists a halfspace h' defining a facet of \mathcal{R}' such that $\mathcal{R}' \subseteq h'$, and h' does not contain x . Since \mathcal{R} intersects $h_{l_1} \cap \dots \cap h_{l_d}$, i) h' intersects $h_{l_1} \cap \dots \cap h_{l_d}$, and ii) h' contains at least an points of P (since $\mathcal{R}' \subseteq h'$), and iii) $|P \cap h' \cap R_{l_i}| \geq bn \forall i \in [1, d]$.

Now, the lowest point z in $\mathcal{R} \cap h'$ is strictly higher than x (since h' does not contain x) and is defined by exactly d halfspaces from \mathcal{H} since \mathcal{R} is full dimensional and is defined by exactly d halfspaces from \mathcal{H} . Furthermore, the set of halfspaces defining z is $\{h'\} \cup \{h_{l_1}, \dots, h_{l_d}\} \setminus h_{l_i}$ for some $i \in [1, d]$ and since $|P \cap C' \cap R_{l_i}| \geq bn \forall i \in [1, d]$, their intersection contains at least bn points from P . This is a contradiction to the assumption that $(h_{l_1}, \dots, h_{l_d})$ has the highest lowest-intersection point among the d -tuples in \mathcal{H}^d . See Figure 7.1 for an example in \mathbb{R}^2 . \square

We now show that any convex set C' containing an points must contain a point of Q by one of the following cases:

1. C' contains x , so is hit by Q .
2. C' does not contain points from \mathcal{R} . Since $|P \cap \mathcal{R}| \geq bn$, C' contains an points from the remaining set $P \setminus (P \cap \mathcal{R})$, whose size is at most $(1 - b)n$. If $an \geq \varepsilon_r^d(1 - b)n$, then C' is hit by Q' .
3. C' does not contain x and yet contains points from \mathcal{R} . Then, by Claim 7.3, $C' \cap R_{l_i} \leq bn$ for some $i \in [1, d]$. Then it must contain at least $an - bn$ points from $P \setminus (P \cap R_{l_i})$. If $an - bn \geq \varepsilon_s^d(1 - ((d - 1)a - (d - 2)))n$, then C' is hit by Q_{l_i} .

Therefore, if

$$an \geq \varepsilon_r^d(1 - b)n \quad \text{and} \quad an - bn \geq \varepsilon_s^d(d - 1)(1 - a)n \quad (7.1)$$

then C' is hit by Q . Maximizing a while satisfying (7.1) yields

$$\varepsilon_{r+ds+1}^d \leq a = \frac{\varepsilon_r^d \cdot (1 + (d - 1)\varepsilon_s^d)}{1 + \varepsilon_r^d \cdot (1 + (d - 1)\varepsilon_s^d)},$$

completing the proof of Lemma 7.2 and hence Theorem 7.1. \square

Remark: The above method actually gives another elementary proof of the centerpoint theorem in any dimension. The proof for two dimensions, as in the method of Theorem 7.1 is: consider all halfspaces containing more than $\frac{2}{3}n$ points, and take the pair with the highest lowest-intersection point x . This is the required point, since any convex set not containing this point cannot intersect the intersection of the halfspaces (Claim 7.3), which contains more than $n/3$ points of P . Hence, such a convex set can only contain the remaining points of P , of which there are fewer than $\frac{2}{3}n$. This follows from Theorem 7.1 by setting $r = s = 0$ and $d = 2$ to get $\varepsilon_1^2 = \frac{2}{3}!$ The proof for d -dimensions is exactly the same: consider sets of d halfspaces, each of which contains more than $\frac{d}{d+1}n$ points and choose the set with the highest lowest-intersection point (w.r.t. any dimension).

7.1 Consequences of main theorem

Improving upon previous work [AAH⁺09], we completely resolve the 2-point case.

Corollary 7.1.1. *Given a set P of n points in \mathbb{R}^2 , the set of all convex objects which contain more than $4n/7$ points of P can be hit by two points (i.e., $\varepsilon_2 \leq 4/7$). Furthermore, there exist arbitrarily large point sets such that the set of all convex objects containing $4n/7$ points cannot be hit by two points.*

Proof. The upper bound follows from Theorem 7.1 by setting $c = 1, d = 0$.

For any n , we construct a point set P of size n such that for any two given points p and q in the plane there is a convex set which avoids both the points and contains $4n/7$ points of P .

Consider the vertices of a regular heptagon each representing a set of $n/7$ points contained in a sufficiently small disk. Let a, b, c, d, e, f and g be the sets in clockwise order. Our set P is the union of these sets.

If one of the points p or q is arbitrarily close to one of the 7 sets, say the set a , then the other point cannot hit the convex hulls of the sets $b \cup c \cup d \cup e, d \cup e \cup f \cup g$ and $f \cup g \cup a \cup b$ simultaneously since they don't have a common intersection. Now, assume that neither p nor q is arbitrarily close to any of the 7 sets. Consider the line l passing through the points p and q . If l does not pass through any of the 7 sets then one of the closed halfspaces defined by l contains 4 of the sets whose convex hull is not hit by either p or q . Otherwise, one of the closed halfspaces defined by l contains 3 of sets and they along with one of the sets which l passes through define a set of $4n/7$ points whose convex hull is not hit by either p or q . \square

Corollary 7.1.2. *Given P , the set of all convex objects which contain more than $8n/15$ points of P can be hit by three points (i.e., $\varepsilon_3 \leq 8/15$). Furthermore, there exist arbitrarily large point sets such that the set of all convex objects containing $5n/11$ points cannot be hit by three points.*

Proof. The upper bound follows from Theorem 7.1 by setting $c = 2, d = 0$. The lower bound construction is as follows.

For any n , we construct a point set P of size n such that for any three given points in the plane there is a convex set containing $5n/11$ points of P which avoids all the three points. Figure 7.2(a) shows such a point set. Each of the 11 points in the figure represents a set of $n/11$ points contained in a sufficiently small disk.

Assume that there are three points which hit all convex sets containing $5n/11$ points of P . We first show that these points cannot be arbitrarily close to any of the 11 sets in the point set. Observe that if all the three points are arbitrarily close to one of the 11 sets in the point set, then they cannot hit the convex region formed by the rest of the 8 sets. Also, if two of the points p, q and r are arbitrarily close to one of the sets, then the rest of the 9 sets can be used to define two convex sets containing $5n/11$ points each and sharing only one of the 11 sets. A single point hitting both these sets should be arbitrarily close to the shared set implying that all the three points are arbitrarily close to one of the sets. If only one of the points is arbitrarily close to a set, say the set k , we take the rest of the 10 sets and consider the convex sets $defgh, fghij$ and $jabcd$. Since two points hit all the three sets,




Figure 7.2: (a) One of the seven (bold) triangles contains a point of the weak ϵ -net (b) One of the four triangles jzk , gxh , dve or asc contains a point of the weak ϵ -net (c) jyi contains a point of the weak ϵ -net

one of the points should be contained in the region $hxfg$. Now, consider the sets $hijab$ and $bcdef$. The third point must hit both these regions and therefore must be arbitrarily close to the set b .

Assuming that none of the points is arbitrarily close to one of the 11 sets, we show that if there exists a set of three points which hits all convex sets containing $5n/11$ points from P then one of those points is contained in one of the bold triangles shown in Figure 7.2(a).

Consider the four convex sets $jkabc$, $abcde$, $defgh$ and $ghijk$ (see Figure 7.2(b)) containing $5n/11$ points each. In order to hit all the four regions, one of the three points must be in one of the four triangles jzk , gxh , dve or asc . If there is a point in one of the triangles jzk , gxh or dve , we are done. So, assume that there is a point in the triangle asc . There cannot be two points in this region since then the remaining one point cannot hit the disjoint regions $ahijk$ and $cdefg$ simultaneously.

If the point in asc is in one of the triangles atb or buc (see Figure 7.2(c)), we are done again. So, we assume that it is in the region $stbu$. But then, the regions $abijk$, $fghij$ and $bcdef$ must be hit by the other two points, and one of those must be in the triangle jyi (see Figure 7.2(c)).

Hence, one of the bold triangles shown in Figure 7.2(a) must contain one of three weak ϵ -net points.

Assume that the triangle hxg contains one of the points (the other cases are analogous). Since the regions $abcdk$, $efijk$ and $defij$ must be hit by two points, the region $efijr$ must contain one of the points (see Figure 7.3(a)). Now, since the regions $abcjk$ and $abcde$ must be hit by one point (see Figure 7.2(a)), the region $abcs$ contains a point.

Also, since the regions $abijk$ and $bcdef$ must be hit (see Figure 7.3(b)), either the regions abt and efw contain one point each or the regions buc and ijy contain one point each. Since the cases are symmetric, let us assume that the regions abt and efw contain one point each.

But then, the region $cdijk$ does not contain any point (see Figure 7.3(c)) although it contains $5n/11$ points of P . Hence, it is not possible to hit all the convex regions containing $5n/11$ points of P using




Figure 7.3: (a) $efijr$ contains a point of the weak ϵ -net (b) Either abt and efw contain one point each or buc and ijy contain one point each. (c) abt , efw and hxg contain one point each. Hence $cdijk$ cannot be hit.

3 points.

□

Aronov [AAH⁺09] show that $\epsilon_4 \leq 4/7$. We actually are able to hit sets containing $4n/7$ points by just two points (Corollary 7.1.1)! For ϵ_4 , Theorem 7.1 yields $16/31$, again improving upon Aronov *et al.*'s result. Improving upon a result of Alon *et al.* [ABFK92], Aronov *et al.* [AAH⁺09] showed that if each convex set contains $n/2$ points, then they can be hit by five points. Theorem 7.1 yields an improvement (set $c = 2$, and $d = 1$).

Corollary 7.1.3. $\epsilon_5 \leq 20/41$.

8. Simplicial Depth

My belt holds my pants up, but the belt loops hold my belt up. I don't really know what's happening down there. Who is the real hero?

Mitch Hedberg

In this chapter we prove the following theorem:

Theorem 8.1. *Any set P of n points in \mathbb{R}^3 has simplicial depth at least $0.0025 \cdot n^4$.*

We remind the reader of the result of Wagner [Wag03], that any point of depth τn is contained in at least the following number of simplices:

$$\frac{(d+1)\tau^d - 2d\tau^{d+1}}{(d+1)!} \cdot n^{d+1} - O(n^d). \quad (8.1)$$

Note that this bound improves with the depth of the point set P . Our simple idea is to show that when the depth of P is low, one can also get a better bound. For example, when $\text{depth}(P) = n/(d+1)$, then the simplicial depth conjecture is in fact proven below. By combining the two approaches, one gets an overall improvement.

We will use the following lemma, which follows easily from a lemma proved in [BF82].

Lemma 8.2. *Given a set P of n points in \mathbb{R}^d , where $\text{depth}(P) = \tau n - 1$, there exists a point p with depth $\tau n - 1$, and a set \mathcal{H} of $d+1$ open halfspaces $\{h_1, \dots, h_{d+1}\}$, such that i) $|h_i \cap P| = \tau n$, ii) p lies on the boundary plane of each h_i , and iii) $h_1 \cup \dots \cup h_{d+1}$ cover the entire \mathbb{R}^d except the point p .*

Proof. Boros-Füredi [BF82] (Lemma 3) prove that given a point set P of size n and depth σ in \mathbb{R}^d , there exists a point p of depth σ and $d+1$ closed halfspaces $\eta_1, \dots, \eta_{d+1}$ which cover \mathbb{R}^d , have the point p on their boundary, and where $|\eta_i \cap P| = n - \sigma - 1$ (they actually prove this statement for \mathbb{R}^2 , but as they also note in their paper, the generalization to \mathbb{R}^d is straightforward). Applying this lemma with $\sigma = \tau n - 1$ and setting h_i to be the complement of η_i for $i = 1, \dots, d+1$, proves the required statement. \square

We now prove a technical lemma which can be seen as a generalization of Carathéodory's theorem. Given a set P of points in \mathbb{R}^d , $\text{conv}(P)$ denotes the convex-hull of P .

Lemma 8.3. *Let $P = \{p_1, \dots, p_{d+2}\}$ be a set of $d + 2$ points in \mathbb{R}^d . Then any point $x \in \text{conv}(P)$ lies in at least two d -simplices spanned by P .*

Proof. If x lies on any facet F of $\text{conv}(P)$, then the (at least) two simplices spanned by P that contain F also contain x . Otherwise take any point of P , say p_1 , and consider the ray emanating from p_1 and passing through x . This ray, after passing through x , intersects the boundary of $\text{conv}(P)$ in a $(d - 1)$ -simplex spanned by d points, say P' . Then $P' \cup p_1$ contains x , and has size $d + 1$. Let p_i be the remaining point in $P \setminus (P' \cup \{p_1\})$. Repeating the same procedure of shooting a ray from p_i through x results in another d -simplex, with p_i as one of its points, that contains x . \square

Given a set P of n points in \mathbb{R}^3 , with $\text{depth}(P) = \tau n - 1$, use Lemma 8.2 to get the point p and a set of four halfspaces $\mathcal{H} = \{h_1, h_2, h_3, h_4\}$ satisfying the stated conditions. The rest of this chapter will be devoted to proving that p is contained in a lot of simplices spanned by P (w.l.o.g. one can assume that $p \notin P$, otherwise the bound can only improve, as all the $\Theta(n^d)$ simplices defined by p contain p). For any halfspace h , let \bar{h} be the complement halfspace of h , and ∂h be its boundary plane. Define the following subsets of P for all $i, j = 1 \dots 4$:

$$A_i = P \cap \left(\bigcap_{l \neq i} \bar{h}_l \right) \cap h_i, \quad B_{i,j} = P \cap \left(\bigcap_{l \neq i,j} \bar{h}_l \right) \cap h_i \cap h_j, \quad C_i = P \cap \left(\bigcap_{l \neq i} h_l \right) \cap \bar{h}_i. \quad (8.2)$$

Set $\alpha_i = |A_i|/n$, $\beta_{i,j} = |B_{i,j}|/n$, and $\gamma_i = |C_i|/n$. Our main lemma is the following.

Lemma 8.4. *Let P be a set of n points in \mathbb{R}^3 , with $\text{depth}(P) = \tau n - 1$. Then, there exists a point contained in at least $g(P) \cdot n^4$ simplices spanned by P , where*

$$g(P) = \left(\prod_i \alpha_i \right) + \left(\sum_{i < j} \beta_{i,j} \cdot \frac{\prod_l \alpha_l}{\max(\alpha_i, \alpha_j)} \right) + \left(\sum_i \gamma_i \cdot \frac{\prod_l \alpha_l}{\max_{l \neq i} \alpha_l} \right). \quad (8.3)$$

Proof. Let p be the point from Lemma 8.2, together with the four halfspaces h_1, \dots, h_4 . We first show that the simplex spanned by any four points, one from each of A_i , will always contain p .

Claim 8.5. *Let $p_1, p_2, p_3, p_4 \in P$ be four points of P , such that $p_i \in A_i$. Then the simplex spanned by these four points contains p .*

Proof. Assume for contradiction that $\text{conv}(\{p_1, p_2, p_3, p_4\})$ does not contain p . Then there exists a hyperplane h that separates p from $\text{conv}(\{p_1, p_2, p_3, p_4\})$, and h does not contain p . For $i \in \{1, 2, 3, 4\}$, define q_i to be the point $pp_i \cap h$. By definition, each h_i passes through p , contains p_i and does not contain any other point p_j with $j \neq i$. Note that then each h_i also contains q_i (by convexity), and does not contain any other point q_j with $j \neq i$. By Radon's theorem [Mat02] applied to $\{q_1, q_2, q_3, q_4\}$ lying on the plane h , there exist disjoint sets Q_1, Q_2 and a point s so that $Q_1 \cup Q_2 = \{q_1, q_2, q_3, q_4\}$ and $s \in \text{conv}(Q_1) \cap \text{conv}(Q_2)$. Since s lies on h , $s \neq p$. By convexity, any halfspace that contains s must also contain at least one point from both Q_1 and Q_2 . As \mathcal{H} covers $\mathbb{R}^3 \setminus \{p\}$, there exists an i such that $s \in h_i$. But this gives a contradiction, as then this h_i must contain at least one point from both Q_1 and Q_2 , and so contain some point q_j with $j \neq i$. \square

The total number of such simplices is $n^4 \cdot \prod_i \alpha_i$, which is the first term in Equation (8.3). Call any such simplex a *basic simplex*, i.e., a simplex on $p_1, p_2, p_3, p_4 \in P$ is basic iff $p_i \in A_i$ for all i . All other simplices are called non-basic.

Now we use basic simplices, which always contain p , to prove the existence of several other simplices which must also contain p .

Claim 8.6. *Let $P' = \{p_1, p_2, p_3, p_4, p_5\} \subset P$ be five points of P , such that $p_k \in A_k$, $k = 1 \dots 4$, and $p_5 \in B_{i,j}$, for any $1 \leq i < j \leq 4$. Then either the simplex spanned by $P' \setminus p_i$ or by $P' \setminus p_j$ contains p .*

Proof. By Claim 8.5, the basic simplex spanned by p_1, p_2, p_3, p_4 contains p . Therefore, by Lemma 8.3, at least one other simplex spanned by P' must contain p . Note that this simplex must have p_5 as one of its points. Also, it must contain p_k , where $k \neq i, j$, since the plane ∂h_k separates $P' \setminus p_k$ from p , as it follows from the definitions (8.2) that $p_l \in \bar{h}_k$ for all $l \neq k$, and $p_5 \in \bar{h}_k$ since $p_5 \in B_{i,j}$. So for this second simplex, the only possible choice is for the fourth vertex, which can be either p_i or p_j . \square

For any fixed i, j , there are $n^5(\beta_{i,j} \cdot \prod_i \alpha_i)$ 5-tuples as in Claim 8.6, and each produces one d -simplex containing p . Each such d -simplex may be double-counted at most $n \cdot \max(\alpha_i, \alpha_j)$ times, so the total number of distinct d -simplices of the type in Claim 8.6 containing p are at least $n^4 \frac{\beta_{i,j} \cdot \prod_i \alpha_i}{\max(\alpha_i, \alpha_j)}$, which when summed over all $i < j$, forms the second term in (8.3).

Claim 8.7. *Let $P' = \{p_1, p_2, p_3, p_4, p_5\} \subset P$ be five points of P , such that $p_k \in A_k$, $k = 1 \dots 4$, and $p_5 \in C_i$, for any $1 \leq i \leq 4$. Then there is a two-element subset $P'' \subset P' \setminus \{p_5, p_i\}$ such that the simplex $\text{conv}(\{p_5, p_i\} \cup P'')$ contains p .*

Proof. As in Claim 8.6, at least one non-basic simplex spanned by P' must contain p , with p_5 as one of its points. Also, it must contain p_i : the plane ∂h_i separates $P' \setminus p_i$ from p , as $P' \setminus p_i \subseteq \bar{h}_i$. The other two vertices of this second simplex must therefore be a subset of the remaining three vertices in P' . \square

By similarly eliminating the double-counting, the d -simplices from Claim 8.7 form the third term of $g(P)$. Finally, note that no two simplices are counted twice in $g(P)$, since each contains exactly one point from a distinct region (one of $B_{i,j}$ or C_i). \square

Note that we only have these two constraints on the non-negative variables $\alpha_i, \beta_{i,j}$ and γ_i :

$$\tau = \frac{|h_i \cap P|}{n} = \alpha_i + \sum_{j \neq i} \beta_{i,j} + \sum_{j \neq i} \gamma_j \quad \text{for each } i = 1 \dots 4. \quad (8.4)$$

$$\sum_i \alpha_i + \sum_{i < j} \beta_{i,j} + \sum_i \gamma_i = 1, \quad \text{as } \mathcal{H} \text{ covers } \mathbb{R}^3 \setminus \{p\} \text{ and } p \notin P. \quad (8.5)$$

It remains to show that regardless of the distribution of the points in the disjoint sets $A_i, B_{i,j}$ and C_j , and therefore the values of the variables satisfying Equations (8.4) and (8.5), the quantity $g(P)$ is always bounded suitably from below.

Lemma 8.8. *Let P be a set of n points in \mathbb{R}^3 , with $\text{depth}(P) = \tau n - 1$, and $g(P)$ as in Lemma 8.4. If $\tau \leq 0.3$, then $g(P) \geq \tau \cdot (1 - 3\tau)^2 \cdot (5\tau - 1)$.*

Proof. Using the fact that $\alpha_i \leq \tau$, we get

$$\begin{aligned}
g(P) &= \left(\prod_i \alpha_i \right) + \left(\sum_{i < j} \beta_{i,j} \cdot \frac{\prod_l \alpha_l}{\max(\alpha_i, \alpha_j)} \right) + \left(\sum_i \gamma_i \cdot \frac{\prod_l \alpha_l}{\max_{l \neq i} \alpha_l} \right) \\
&\geq \left(\prod_i \alpha_i \right) + \left(\sum_{i < j} \beta_{i,j} \cdot \frac{\prod_l \alpha_l}{\tau} \right) + \left(\sum_i \gamma_i \cdot \frac{\prod_l \alpha_l}{\tau} \right) \\
&= \left(\prod_i \alpha_i \right) \left(1 + \frac{\sum_{i < j} \beta_{i,j} + \sum_i \gamma_i}{\tau} \right). \tag{8.6}
\end{aligned}$$

Summing up (8.4) for all four halfspaces, and subtracting (8.5) from it, we get

$$\sum_{i < j} \beta_{i,j} + 2 \cdot \sum_i \gamma_i = 4 \cdot \tau - 1. \tag{8.7}$$

Therefore, $\sum_{i < j} \beta_{i,j} + \sum_i \gamma_i \leq 4\tau - 1$. This fact, together with equation (8.4), implies that $1 - 3\tau \leq \alpha_i \leq \tau$ for $i = 1 \dots 4$. Assuming $\tau \leq 0.3$, we can show the following:

Claim 8.9. *The bound in equation (8.6) is minimized when $\sum_i \gamma_i = 0$ or equivalently, when $\sum_{i < j} \beta_{i,j} + \sum_i \gamma_i = 4\tau - 1$.*

Proof. Suppose that $\sum_i \gamma_i = \epsilon + \epsilon_1$, where $\gamma_1 = \epsilon_1 > 0$. We show that the variables $\alpha_i, \beta_{i,j}, \gamma_1$ can be re-adjusted to new values $\alpha'_i, \beta'_{i,j}, \gamma'_1$ to give a smaller value in Equation (8.6), while still satisfying all the constraints in Equations (8.4) and (8.5), and where $\gamma'_1 = 0$. As long as $\sum_i \gamma_i > 0$, we can iteratively apply this procedure for all $\gamma_j > 0$ to make $\sum_i \gamma_i = 0$ without increasing the lower bound. At each such step, a value of $\gamma_j > 0$, for some j , is set to 0, so this procedure finishes after at most 4 steps.

Set $\gamma'_1 = 0$, $\beta'_{i,j} = \beta_{i,j} + \frac{2\epsilon_1}{3}$ for all $i, j \neq 1$, and $\alpha'_i = \alpha_i - \epsilon_1/3$ for all $i \neq 1$. One can verify that Equations (8.4) and (8.5) still hold. Therefore, Equation (8.7) also holds; in particular, it follows that $1 - 3\tau \leq \alpha'_i \leq \tau$ for each i , and so all new α'_i variables are still non-negative as $\tau \leq 0.3$. Now simple calculation shows that the function of Equation (8.6) can only decrease. For completeness sake, we present the explicit computations:

$$\left(\prod_i \alpha'_i \right) \left(1 + \frac{\sum \gamma'_i + \sum \beta'_{i,j}}{\tau} \right) \leq \left(\prod_i \alpha_i \right) \left(1 + \frac{\sum \gamma_i + \sum \beta_{i,j}}{\tau} \right).$$

Note that $\sum_{i < j} \beta_{i,j} + \sum_i \gamma_i = 4\tau - 1 - \epsilon - \epsilon_1$, and $\sum_{i < j} \beta'_{i,j} + \sum_i \gamma'_i = 4\tau - 1 - \epsilon$. So

$$\begin{aligned} \left(\prod_{i \neq 1} \alpha_i - \frac{\epsilon_1}{3} \right) \left(\frac{5\tau - 1 - \epsilon}{\tau} \right) &\leq \left(\prod_{i \neq 1} \alpha_i \right) \left(\frac{5\tau - 1 - \epsilon - \epsilon_1}{\tau} \right) \\ \prod_{i \neq 1} \left(1 - \frac{\epsilon_1}{3\alpha_i} \right) &\leq 1 - \frac{\epsilon_1}{5\tau - 1 - \epsilon} \end{aligned}$$

Since each α_i , $i \neq 1$, can be at most $\tau - \epsilon_1$ (from (8.4)), we have to prove

$$\left(1 - \frac{\epsilon_1}{3(\tau - \epsilon_1)} \right)^3 = 1 - \frac{3\epsilon_1}{3(\tau - \epsilon_1)} + 3 \left(\frac{\epsilon_1}{3(\tau - \epsilon_1)} \right)^2 - \left(\frac{\epsilon_1}{3(\tau - \epsilon_1)} \right)^3 \leq 1 - \frac{\epsilon_1}{5\tau - 1 - \epsilon}$$

By dropping the negative cubic term and simplifying, we have to show

$$\begin{aligned} 1 - \frac{3\epsilon_1}{3(\tau - \epsilon_1)} + 3 \left(\frac{\epsilon_1}{3(\tau - \epsilon_1)} \right)^2 &\leq 1 - \frac{\epsilon_1}{5\tau - 1 - \epsilon} \\ -\frac{1}{(\tau - \epsilon_1)} + \frac{\epsilon_1}{3(\tau - \epsilon_1)^2} &\leq -\frac{1}{5\tau - 1 - \epsilon} \\ -(5\tau - 1 - \epsilon) + \frac{\epsilon_1(5\tau - 1 - \epsilon)}{3(\tau - \epsilon_1)} &\leq -(\tau - \epsilon_1) \\ 4\tau &\geq 1 + \epsilon - \epsilon_1 + \frac{\epsilon_1(5\tau - 1 - \epsilon)}{3(\tau - \epsilon_1)}. \end{aligned}$$

Since $\sum_{i < j} \beta_{i,j} + 2 \sum_i \gamma_i = 4\tau - 1$, we have $\sum \gamma_i = \epsilon + \epsilon_1 \leq (4\tau - 1)/2$. So $4\tau \geq 2\epsilon + 1$, and it remains to show that

$$\frac{\epsilon_1(5\tau - 1 - \epsilon)}{3(\tau - \epsilon_1)} \leq \epsilon_1 \quad \text{or equivalently,} \quad 2\tau \leq 1 + \epsilon - 3\epsilon_1.$$

Now one can verify that $2\tau \leq 1 - 3\epsilon_1$, since $\epsilon_1 \leq (4\tau - 1)/2$, and $\tau \leq 0.3$. □

It follows from Claim 8.9 that

$$g(P) \geq \left(\prod_i \alpha_i \right) \left(1 + \frac{4\tau - 1}{\tau} \right) = \left(\prod_i \alpha_i \right) \left(\frac{5\tau - 1}{\tau} \right).$$

Claim 8.9 together with (8.5) also implies that $\sum \alpha_i = 2 - 4\tau$. As $\alpha_i \in [1 - 3\tau, \tau]$, the term $\prod \alpha_i$ is minimized when, say, $\alpha_1 = \alpha_2 = \tau$, and $\alpha_3 = \alpha_4 = 1 - 3\tau$ (and then $\beta_{3,4} = 4\tau - 1$).

Claim 8.10. $\prod \alpha_i$ is minimized when $\alpha_1 = \alpha_2 = \tau$, and $\alpha_3 = \alpha_4 = (1 - 3\tau)$.

Proof. Recall that each α_i lies in the closed interval $[(1 - 3\tau), \tau]$. If each $\alpha_i > (1 - 3\tau)$, pick the smallest of them, say α_4 , and set $\alpha'_4 = \alpha_4 - \varepsilon$, for a small enough $\varepsilon > 0$, and add this excess to any other variable that is less than τ , say α_3 (there always exists another variable less than τ , else $\sum \alpha_i > (1 - 3\tau) + 3\tau = 1$, a contradiction). Then $(\alpha_3 + \varepsilon)(\alpha_4 - \varepsilon) < \alpha_3\alpha_4$ since $\alpha_4 \leq \alpha_3$, minimizing the product further. Similarly, α_3 is also $(1 - 3\tau)$ in the configuration minimizing $\prod \alpha_i$. So we get that $\alpha_1 + \alpha_2 = (2 - 4\tau) - 2(1 - 3\tau) = 2\tau$. And as each α is at most τ , this forces $\alpha_1 = \alpha_2 = \tau$. \square

It can be verified that all the constraints are satisfied, and so we get the required lower bound for $g(P)$:

$$g(P) \geq \left(\prod_i \alpha_i \right) \left(\frac{5\tau - 1}{\tau} \right) \geq \tau^2 \cdot (1 - 3\tau)^2 \cdot \frac{5\tau - 1}{\tau}.$$

\square

We can now complete the proof of Theorem 8.1. Wagner [Wag03] proved that any point of depth $\tau \cdot n$ in \mathbb{R}^3 is contained in at least $f(\tau) \cdot n^4$ simplices spanned by P , where $f(\tau) = (4\tau^3 - 6\tau^4)/4!$ and $0.25 \leq \tau \leq 0.5$. If P has depth at least $\tau \cdot n$, where $\tau \geq 0.2889$, then as $f'(\tau) \geq 0$ for $\tau \in [0.25, 0.5]$, we can deduce that $f(\tau) \geq f(0.2889) = 0.00227$ and so there exists a point lying in at least $0.00227 \cdot n^4$ simplices spanned by P .

Otherwise, as depth is always an integer, P has depth at most $\tau \cdot n - 1$, where $\tau \leq 0.2889$. By Lemma 8.8, we can conclude that there exists a point lying in $g(\tau) \cdot n^4$ simplices, where $g(\tau) = \tau \cdot (1 - 3\tau)^2 \cdot (5\tau - 1)$. As $g'(\tau) \leq 0$ for $\tau \in [0.25, 0.3]$, we can deduce that $g(\tau) \geq g(0.2889) = 0.00227$ and so there exists a point lying in at least $0.00227 \cdot n^4$ simplices spanned by P .

9. Oja Depth

Meeting a friend in the corridor, Ludwig Wittgenstein said: “Tell me, why do people always say that it was natural for men to assume that the sun went around the earth rather than the earth was rotating?”

His friend said: “Well, obviously, because it just looks as if the sun is going around the earth.” To which the philosopher replied: “Well, what would it look like if it had looked as if the earth were rotating?”

In this chapter we prove the optimal bound in \mathbb{R}^2 and improve the previous-best bound for \mathbb{R}^3 . In Section 9.1, we resolve the planar case of the Oja-depth conjecture. In particular we prove that for every set P of n points in \mathbb{R}^2 the center of mass of the convex hull of P has depth at most $\frac{n^2}{9}$.

Theorem 9.1. *Every set P of n points in \mathbb{R}^2 has Oja depth at most $\frac{n^2}{9}$. Furthermore, this depth is attained by the centroid of the point set which be computed in $O(n \log n)$ time.*

In Section 9.2, using completely different (and more combinatorial) techniques for higher dimensions, we prove:

Theorem 9.2. *Every set P of n points in \mathbb{R}^d , $d \geq 3$, has Oja depth at most*

$$\frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2 (d+1)!} \binom{n}{d} + O(n^{d-1}).$$

9.1 A tight bound in \mathbb{R}^2

We now come to prove a tight bound for \mathbb{R}^2 . The *center of mass* or *centroid* of a convex set X is defined as

$$c(X) = \frac{\int_{x \in X} x \, dx}{\text{area}(X)}.$$

For a discrete point set P , the center of mass of P is defined as the center of mass of the convex hull of P . When we talk about the *centroid of P* , we refer to the center of mass of the convex hull (not to be confused with the discrete centroid $\sum p/|P|$). We will bound the Oja depth of the centroid of a set. As we will see the Oja depth of the centroid is tight in the worst case. Our proof will rely on the following two known results.

Lemma 9.3. [Winternitz [Bla23]] Every line through the centroid of a convex object has at most $\frac{5}{9}$ of the total area on either side.

Lemma 9.4. [CDI⁺13] Let P be a convex object with unit area and let c be its centroid. Then every simplex inside P that has c as a vertex has area at most $\frac{1}{3}$.

To simplify matters, we will use the following proposition.

Proposition 1. If we project an interior point $p \in P$ radially outwards from the centroid c to the boundary of the convex hull, the Oja depth of the point c does not decrease.

Proof. First, observe that the center of mass does not change. It suffices to show that every triangle that has p as one of its vertices increases its area. Let $T := \Delta(c, p, q)$ be any triangle. The area of T is $\frac{1}{2}\|c - p\| \cdot h$, where h is the height of T with respect to $p - c$. If we move p radially outwards to a point p' , h does not change, but $\|c - p'\| > \|c - p\|$. See Fig. 9.1. \square




Figure 9.1: Moving points to the boundary increases the Oja depth

This implies that in order to prove an upper bound, we can assume that P is in convex position. Note that the aforementioned transformation brings the point only in weakly convex position, that is, some of the points lying on the boundary of the convex hull might not actually be vertices of the convex hull. This, however, is sufficient for our proof and for brevity we will use “convex” to mean “weakly convex”.

From now on, let P be a set of points in convex position, and let $c := c(\text{conv}(P))$ denote its center of mass as defined above. Further, let p_1, \dots, p_n denote the points sorted clockwise by angle from c . We define the *distance* of two points p_i, p_j , $i \neq j$, as the difference of their position in this order

(modulo n):

$$\text{dist}(p_i, p_j) := \min\{j - i \pmod n, i - j \pmod n\} \in \{1, \dots, \lfloor n/2 \rfloor\}.$$

A triangle that is formed by c and two points at distance i is called an i -triangle, or triangle of type i . Observe that for each i , $1 \leq i < \lfloor n/2 \rfloor$, there are exactly n triangles of type i . Further, if n is even, then there are $n/2$ triangles of type $\lfloor n/2 \rfloor$, otherwise there are n . These constitute all possible triangles.

Let $C \subseteq P$, and let C be the boundary of the convex hull of C . This polygon will be called a cycle. The length of a cycle is simply the number of elements in C . A cycle C of length i induces i triangles that arise by taking all triangles formed by an edge in C and the center of mass c (of $\text{conv}(P)$). The area induced by C is the sum of areas of these i triangles. See Fig. 9.1.




Figure 9.2: Cycles

The triangles induced by $C = P$ form a partition of $\text{conv}(P)$. Thus

Lemma 9.5. *The total area of all triangles of type 1 is exactly 1.*

The following shows that we can generalize this lemma to bound the total area induced by any cycle:

Lemma 9.6. *Let C be a cycle. Then C induces a total area of at most 1.*

Proof. We distinguish two cases.

Case 1: The centroid c lies inside C . In this case, all triangles are disjoint, so the area is at most one.

Case 2: c does not lie inside C (see Fig. 9.1). Then there is a line through c that has all the triangles induced by C on one side. Then we can remove one triangle – the one induced by the pair $\{p_{i_j}, p_{i_{j+1}}\}$ that has c on the left side (the gray triangle shown in the figure) – to get a set of disjoint triangles. By Lemma 9.3, the area of the remaining triangles can thus be at most $5/9$. By Lemma 9.4, the removed triangle has an area of at most $1/3$. Thus, the total area is at most $8/9$.

□

We now prove the key lemma, which is a general version of Lemma 9.5.

Lemma 9.7. *The total area of all triangles of type i is at most i .*

Proof. To prove this lemma for a fixed i , we will create n cycles. Each cycle will consist of one triangle of type i , and $n - i$ triangles of type 1 (counting multiplicities). We then determine the total area of these cycles and subtract the area of all 1-triangles. This will give the desired result.

Let p_1, \dots, p_n be the points ordered by angles from the centroid c (recall that we could assume that P is in convex position). For $j = 1 \dots n$, let \mathcal{C}_j be the cycle consisting of the $n - i + 1$ points of P $\{p_{j+i \bmod n}, p_{j+i+1 \bmod n}, \dots, p_{j-1 \bmod n}, p_j \bmod n\}$. This is a cycle that consists of one triangle of type i (the one defined by the three points c, p_j, p_{j+i}), and $n - i$ triangles of type 1.

By Lemma 9.6, every cycle \mathcal{C}_j induces an area of at most 1. If we sum up the areas of all cycles \mathcal{C}_j , $1 \leq j \leq n$, we thus get an area of at most n .

We now determine how often we have counted each triangle. Each i -triangle is counted exactly once. Further, for every cycle we count $n - i$ triangles of type 1. For reasons of symmetry, each 1-triangle is counted equally often. Indeed, each one is counted *exactly* $n - i$ times over all the cycles. By Lemma 9.5, their area is *exactly* $n - i$, which we can subtract from n to get the total area of the i -triangles:

$$\sum_{i\text{-triangle } T} \text{area}(T) \leq n - (n - i) \cdot \left(\sum_{1\text{-triangle } T} \text{area}(T) \right) = n - (n - i) = i.$$

□

Now we prove the main result of this section:

Theorem 9.8. *Let P be any set of points in the plane with $\text{area}(\text{conv}(P)) = 1$, and let c be the centroid of $\text{conv}(P)$. Then the Oja depth of c is at most $\frac{n^2}{9}$.*

Proof. We will bound the area of the triangles depending on their type. For i -triangles with $1 \leq i \leq \lfloor n/3 \rfloor$, we will use Lemma 9.7. For i -triangles with $\lfloor n/3 \rfloor < i \leq \lfloor n/2 \rfloor$, this would give us a bound worse than $n/3$, so we will use Lemma 9.4 for each of these.

By Lemma 9.7, the sum of the areas of all triangles of type at most $\lfloor n/3 \rfloor$ is at most

$$\sum_{i=1}^{\lfloor n/3 \rfloor} i = \frac{\lfloor n/3 \rfloor (\lfloor n/3 \rfloor + 1)}{2} = \frac{1}{2} \lfloor \frac{n}{3} \rfloor^2 + \frac{1}{2} \lfloor \frac{n}{3} \rfloor.$$

If n is odd, there are $n (\lfloor n/2 \rfloor - \lfloor n/3 \rfloor)$ triangles remaining, n for each type j , $\lfloor n/3 \rfloor < j \leq \lfloor n/2 \rfloor$. If n is even, there are only $n/2$ triangles of type $n/2$ and so $n (\lfloor n/2 \rfloor - \lfloor n/3 \rfloor - 1/2)$ triangles remaining. In either case the number of remaining triangles is $n^2/2 - n \lfloor n/3 \rfloor - n/2$. For these we use Lemma 9.4 to bound the area of each by $1/3$. Thus, the area of these remaining triangles is at most $\frac{n^2}{6} - \frac{n}{3} \lfloor \frac{n}{3} \rfloor - \frac{n}{6}$.

So the Oja depth is at most $\frac{1}{2} \lfloor \frac{n}{3} \rfloor^2 + \frac{1}{2} \lfloor \frac{n}{3} \rfloor + \frac{n^2}{6} - \frac{n}{3} \lfloor \frac{n}{3} \rfloor - \frac{n}{6}$.

To complete the proof we distinguish the cases when n is of the form $3k$, $3k + 1$ or $3k + 2$.

Case $n = 3k$:

$$\frac{1}{2} \lfloor \frac{n}{3} \rfloor^2 + \frac{1}{2} \lfloor \frac{n}{3} \rfloor + \frac{n^2}{6} - \frac{n}{3} \lfloor \frac{n}{3} \rfloor - \frac{n}{6} = \frac{k^2}{2} + \frac{k}{2} + \frac{3k^2}{2} - k^2 - \frac{k}{2} = k^2 = \frac{n^2}{9}$$

Case $n = 3k + 1$:

$$\frac{1}{2} \lfloor \frac{n}{3} \rfloor^2 + \frac{1}{2} \lfloor \frac{n}{3} \rfloor + \frac{n^2}{6} - \frac{n}{3} \lfloor \frac{n}{3} \rfloor - \frac{n}{6} = \frac{k^2}{2} + \frac{k}{2} + \frac{3k^2}{2} + k + \frac{1}{6} - k^2 - \frac{k}{3} - \frac{k}{2} - \frac{1}{6} = k^2 + \frac{2k}{3} \leq \frac{(3k+1)^2}{9} = \frac{n^2}{9}$$

Case $n = 3k + 2$:

$$\frac{1}{2} \lfloor \frac{n}{3} \rfloor^2 + \frac{1}{2} \lfloor \frac{n}{3} \rfloor + \frac{n^2}{6} - \frac{n}{3} \lfloor \frac{n}{3} \rfloor - \frac{n}{6} = \frac{k^2}{2} + \frac{k}{2} + \frac{3k^2}{2} + 2k + \frac{2}{3} - k^2 - \frac{2k}{3} - \frac{k}{2} - \frac{1}{3} = k^2 + \frac{4k}{3} + \frac{1}{3} \leq \frac{(3k+2)^2}{9} = \frac{n^2}{9}$$

Thus, the Oja depth of the centroid is at most $n^2/9$. □

Remark: We note that $c = c(\text{conv}(P))$ can be computed in $O(n \log n)$ time by first computing $\text{conv}(P)$ in $O(n \log n)$ time, and then triangulating $\text{conv}(P)$.

9.2 Higher Dimensions

We now present improved bounds for the Oja depth problem in dimensions greater than two. Before the main theorem, we need the following two lemmas.

Lemma 9.9. *Given any set P of n points in \mathbb{R}^d and any point $q \in \mathbb{R}^d$, any line l through q intersects at most $f(n, d)$ $(d - 1)$ -simplices spanned by P , where $f(n, d) = \frac{2n^d}{2^{d-1}d!} + O(n^{d-1})$.*

Proof. Project P onto the hyperplane H orthogonal to l to get the point set P' in \mathbb{R}^{d-1} . The line l becomes a point on H , say point l^* . Then l intersects the $(d-1)$ -simplex spanned by $\{p_1, \dots, p_d\}$ if and only if the convex hull of the corresponding points in P' contains the point l^* . By Bárány [Bár82], given n points in \mathbb{R}^k , any point in \mathbb{R}^k is contained in at most this many k -simplices:

$$\begin{cases} \frac{2n}{n+k+1} \cdot \binom{(n+k+1)/2}{k+1} & \text{if } n-k \text{ is odd} \\ \frac{2(n-k)}{n+k+2} \cdot \binom{(n+k+2)/2}{k+1} & \text{if } n-k \text{ is even} \end{cases}$$

Note that both the bounds above are within an additive term of $O(n^k)$, and simplifying the first, we get:

$$\begin{aligned} \frac{2n}{n+k+1} \cdot \binom{(n+k+1)/2}{k+1} &\leq 2 \cdot \binom{(n+k+1)/2}{k+1} \\ &\leq \frac{2(n+k+1)^{k+1}}{2^{k+1}(k+1)!} \leq \frac{2n^{k+1}}{2^{k+1}(k+1)!} + O(n^k) \end{aligned}$$

We apply this to P' in \mathbb{R}^{d-1} (setting k to $d - 1$) to get the desired result. □

Lemma 9.10. *Given any set P of n points in \mathbb{R}^d , there exists a point q such that any half-infinite ray from q intersects at least $\frac{2^d}{(d+1)^2(d+1)!} \binom{n}{d}$ $(d-1)$ -simplices spanned by P .*

Proof. This follows directly from a recent result of Gromov [Gro10], who showed that given any set P , there exists a point q contained in at least $\frac{2^d}{(d+1)(d+1)!} \binom{n}{d+1}$ simplices spanned by P . Now note that any half-infinite ray from q must intersect exactly one $(d-1)$ -dimensional face of each simplex containing q and each such $(d-1)$ -simplex can be counted at most $n-d$ times. Simplifying, we get the desired result. \square

Remark. There have been several improvements [MW10, KMS12] after the initial work of Gromov; however as these improvements are significant only for small constant dimensions, we prefer to give the above considerably simpler bound of Gromov. It is clear that any improvement in the above bound gives a corresponding improvement for our result.

Given a set P and a point q , call a simplex a q -simplex if it is spanned by q and d other points of P .

Theorem 9.11. *Given any set P of n points in \mathbb{R}^d , there exists a point q with Oja depth at most*

$$\frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2(d+1)!} \binom{n}{d} + O(n^{d-1}).$$

Proof. Let q be the point from Lemma 9.10. Assign a weight function, $w(r)$, to each point $r \in \text{conv}(P)$, where $w(r)$ is the number of q -simplices spanned by P and q that contain r . Then note that if r is contained in a q -simplex, spanned by, say, $\{q, p_{i_1}, \dots, p_{i_d}\}$, then the half-infinite ray \vec{qr} intersects the $(d-1)$ -simplex spanned by $\{p_{i_1}, \dots, p_{i_d}\}$. Therefore $w(r)$ is equal to the number of $(d-1)$ -simplices intersected by the ray \vec{qr} . To upper-bound this, note that the ray starting from q but in the opposite direction to the ray \vec{qr} , intersects at least $\frac{2^d}{(d+1)^2(d+1)!} \binom{n}{d}$ $(d-1)$ -simplices (by Lemma 9.10). On the other hand, by Lemma 9.9, the entire line passing through q and r intersects at most $\frac{2n^d}{2^d d!} + O(n^{d-1})$ $(d-1)$ -simplices spanned by P . These two together imply that the ray \vec{qr} intersects at most $\frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2(d+1)!} \binom{n}{d} + O(n^{d-1})$ $(d-1)$ -simplices spanned by P , and this is also an upper bound on $w(r)$. Finally, we have

$$\begin{aligned} \text{Oja-depth}(q, P) &= \sum_{P' \subseteq P, |P'|=d} \text{vol}(\text{conv}(\{q\} \cup P')) = \int_{x \in \text{conv}(P)} w(x) dx \\ &\leq \left(\frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2(d+1)!} \binom{n}{d} + O(n^{d-1}) \right) \int_{x \in \text{conv}(P)} dx \\ &= \frac{2n^d}{2^d d!} - \frac{2d}{(d+1)^2(d+1)!} \binom{n}{d} + O(n^{d-1}), \end{aligned}$$

finishing the proof. \square

Appendices

A. List of Basic Theorems

For a set $P \subset \mathbb{R}^d$, $\text{conv}(P)$ denotes the convex hull of P .

Carathéodory's Theorem. If a point $q \in \mathbb{R}^d$ lies in the convex-hull of a set of points P , then it also lies in the convex-hull of a set $Q \subset P$ of size at most $d + 1$.

Colorful Carathéodory's Theorem. Given $d + 1$ sets of points P_1, \dots, P_{d+1} in \mathbb{R}^d and a point q such that $q \in \text{conv}(P_i)$ for all $i = 1, \dots, d + 1$, there exists a set Q such that $q \in \text{conv}(Q)$ and where $|Q \cap P_i| = 1$ for all i . Such a Q is called a 'rainbow set'.

Helly's Theorem. Given a set \mathcal{C} of compact convex objects in \mathbb{R}^d such that every $(d + 1)$ of them have a common intersection, all of them have a common intersection.

VC-dimension. Given a range space (X, \mathcal{F}) , a set $X' \subseteq X$ is *shattered* if every subset of X' can be obtained by intersecting X' with a member of the family \mathcal{F} . The VC-dimension of (X, \mathcal{F}) is the size of the largest set that can be shattered.

ε -nets. The ε -net theorem [HW87] states that there exists a ε -net of size $O(d/\varepsilon \log 1/\varepsilon)$ for any range space with VC-dimension d . This bound was later improved in [KPW92] to a near-optimal bound of $(1 + o(1))(\frac{d}{\varepsilon} \log(1/\varepsilon))$. For example, given a set P of n points in \mathbb{R}^d and a parameter $\varepsilon > 0$, a set $Q \subseteq P$ is an ε -net w.r.t. halfspaces if any halfspace containing at least εn points of P contains a point of Q .

Radon's theorem. Any set of $d + 2$ points in \mathbb{R}^d can be partitioned into two sets A and B such that $\text{conv}(A) \cap \text{conv}(B) \neq \emptyset$.

Tverberg's theorem. Given a set P of $(r - 2)(d + 1) + (d + 2)$ points in \mathbb{R}^d , one can partition P into r sets P_1, \dots, P_r , such that there exists a point p lying in each $\text{conv}(P_i)$, i.e., $\bigcap_i \text{conv}(P_i) \neq \emptyset$.

Kirchberger's Theorem. Let P_1, \dots, P_r be r sets of points in \mathbb{R}^d whose convex-hulls have non-empty intersection, i.e., $\cap_i \text{conv}(P_i) \neq \emptyset$. Then there exist subsets $P'_1 \subseteq P_1, \dots, P'_r \subseteq P_r$ of total size $(r-1)(d+1)+1$ which also have a non-empty common intersection, i.e., $\cap_i \text{conv}(P'_i) \neq \emptyset$, and $|\cup_i P'_i| = (r-1)(d+1)+1$.

Ramsey's theorem for hypergraphs. There exists a function $f(n)$ such that given any 2-coloring of the edges of a complete k -uniform hypergraph on at least $f(n)$ vertices, there exists a subset of size n such that all edges induced by this subset are monochromatic.

Hadwiger-Debrunner (p, q) -theorem. Given a set S of convex objects in \mathbb{R}^d such that out of every $p \geq d+1$ objects, there is a point common to $q \geq (d+1)$ of them, then S has a hitting set of size $HD_d(p, q)$ (independent of $|S|$).

Centerpoint depth. Given any set P of n points in \mathcal{R}^d , the *Tukey depth* of a point $q \in \mathbb{R}^d$ is the minimum number of points of P contained in any halfspace containing q . It is known that there always exists a point of Tukey depth at least $n/(d+1)$.

PUBLICATIONS CO-AUTHORED BY THE CANDIDATE

- [1] Pankaj K. Agarwal, Sarel Har-Peled, Nabil H. Mustafa, and Yusu Wang. Near-linear time approximation algorithms for curve simplification. In *European Symposium on Algorithms (ESA)*, pages 29–41, 2002.
- [2] Pankaj K. Agarwal, Sarel Har-Peled, Nabil H. Mustafa, and Yusu Wang. Near-linear time approximation algorithms for curve simplification. *Algorithmica*, 42(3-4):203–219, 2005.
- [3] Pankaj K. Agarwal, Shankar Krishnan, Nabil H. Mustafa, and Suresh Venkatasubramanian. Streaming geometric optimization using graphics hardware. In *European Symposium on Algorithms (ESA)*, pages 544–555, 2003.
- [4] Pankaj K. Agarwal and Nabil H. Mustafa. Independent set of intersection graphs of convex objects in 2D. In *Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, pages 127–137, 2004.
- [5] Pankaj K. Agarwal and Nabil H. Mustafa. k -means projective clustering. In *ACM Symposium on Principles of Database Systems (PODS)*, pages 155–165, 2004.
- [6] Pankaj K. Agarwal and Nabil H. Mustafa. Independent set of intersection graphs of convex objects in 2D. *Comput. Geom.*, 34(2):83–95, 2006.
- [7] Pankaj K. Agarwal, Nabil H. Mustafa, and Yusu Wang. Fast molecular shape matching using contact maps. *Journal of Computational Biology*, 14(2):131–143, 2007.
- [8] Yih-En Andrew Ban, Sergey Bereg, and Nabil H. Mustafa. A conjecture on Wiener indices in combinatorial chemistry. *Algorithmica*, 40(2):99–117, 2004.
- [9] Yih-En Andrew Ban, Sergei Bespamyatnikh, and Nabil H. Mustafa. On a conjecture on Wiener indices in combinatorial chemistry. In *International Computing and Combinatorics Conference (COCOON)*, pages 509–518, 2003.
- [10] Abdul Basit, Nabil H. Mustafa, Saurabh Ray, and Sarfraz Raza. Centerpoints and Tverberg’s technique. *Comput. Geom.*, 43(6-7):593–600, 2010.
- [11] Abdul Basit, Nabil H. Mustafa, Saurabh Ray, and Sarfraz Raza. Hitting simplices with points in \mathbb{R}^3 . *Discrete & Computational Geometry*, 44(3):637–644, 2010.
- [12] Abdul Basit, Nabil H. Mustafa, Saurabh Ray, and Sarfraz Raza. Improving the First Selection lemma in \mathbb{R}^3 . In *Symposium on Computational Geometry*, pages 354–357, 2010.
- [13] Khaled M. Elbassioni, Aleksei V. Fishkin, Nabil H. Mustafa, and René Sitters. Approximation algorithms for Euclidean group TSP. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 1115–1126, 2005.
- [14] Khaled M. Elbassioni and Nabil H. Mustafa. Conflict-free colorings of rectangles ranges. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 254–263, 2006.
- [15] Shankar Krishnan, Nabil H. Mustafa, and Suresh Venkatasubramanian. Hardware-assisted computation of depth contours. In *ACM - SIAM Symposium on Discrete Algorithms (SODA)*, pages 558–567, 2002.
- [16] Shankar Krishnan, Nabil H. Mustafa, and Suresh Venkatasubramanian. Statistical Data Depth and the Graphics Hardware. In *Data Depth: Robust Multivariate Analysis, Computational Geometry and Applications*, pages 223–250. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 2006.

- [17] Nabil H. Mustafa. Simplification of geometric objects. Master's thesis, Duke University, Durham, North Carolina, 2002.
- [18] Nabil H. Mustafa. *Simplification, Estimation and Classification of Geometric Objects*. PhD thesis, Duke University, Durham, North Carolina, 2004.
- [19] Nabil H. Mustafa, Eleftherios Koutsofios, Shankar Krishnan, and Suresh Venkatasubramanian. Hardware-assisted view-dependent map simplification. In *Symposium on Computational Geometry*, pages 50–59, 2001.
- [20] Nabil H. Mustafa, Shankar Krishnan, Gokul Varadhan, and Suresh Venkatasubramanian. Dynamic simplification and visualization of large maps. *International Journal of Geographical Information Science*, 20(3):273–302, 2006.
- [21] Nabil H. Mustafa and Aleksandar Pekec. Majority consensus and the local majority rule. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 530–542, 2001.
- [22] Nabil H. Mustafa and Aleksandar Pekec. Listen to your neighbors: How (not) to reach a consensus. *SIAM J. Discrete Math.*, 17(4):634–660, 2004.
- [23] Nabil H. Mustafa and Saurabh Ray. Weak ϵ -nets have a basis of size $O(1/\epsilon \log 1/\epsilon)$. *Comp. Geom: Theory and Appl.*, 40(1):84–91, 2008.
- [24] Nabil H. Mustafa and Saurabh Ray. Near-optimal generalizations of a theorem of Macbeath. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, 2014.
- [25] Nabil H. Mustafa and Saurabh Ray. An optimal generalization of the centerpoint theorem, and its extensions. In *Symposium on Computational Geometry*, pages 138–141, 2007.
- [26] Nabil H. Mustafa and Saurabh Ray. Weak ϵ -nets have basis of size $O(1/\epsilon \log(1/\epsilon))$ in any dimension. In *Symposium on Computational Geometry*, pages 239–244, 2007.
- [27] Nabil H. Mustafa and Saurabh Ray. An optimal extension of the Centerpoint theorem. *Comput. Geom.*, 42(6-7):505–510, 2009.
- [28] Nabil H. Mustafa and Saurabh Ray. PTAS for geometric hitting set problems via local search. In *Symposium on Computational Geometry*, pages 17–22, 2009.
- [29] Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010.
- [30] Nabil H. Mustafa and Saurabh Ray. A theorem of Bárány revisited and extended. In *Symposium on Computational Geometry*, pages 333–338, 2012.
- [31] Nabil H. Mustafa, Saurabh Ray, and Mudassir Shabbir. Ray-shooting depth: Computing statistical data depth of point sets in the plane. In *European Symposium on Algorithms (ESA)*, pages 506–517, 2011.
- [32] Nabil H. Mustafa, Saurabh Ray, and M. Shabbir. Centerflats: A unifying view of data-depth. *Submitted*, 2013.
- [33] Nabil H. Mustafa, H. Tiwary, and D. Werner. A proof of the Oja-depth conjecture in the plane. *Comp. Geom: Theory and Appl.*, 2013.

OTHER REFERENCES

- [AAH⁺09] Boris Aronov, Franz Aurenhammer, Ferran Hurtado, Stefan Langerman, David Rappaport, Carlos Seara, and Shakhar Smorodinsky. Small weak epsilon-nets. *Comput. Geom.*, 42(5):455–462, 2009.
- [ABET00] N. Amenta, M. Bern, D. Eppstein, and S. Teng. Regression depth and center points. *Discrete & Computational Geometry*, 23(3):305–323, 2000.
- [ABFK92] Noga Alon, Imre Bárány, Zoltán Füredi, and Daniel J. Kleitman. Point selections and weak ϵ -nets for convex hulls. *Combinatorics, Probability & Computing*, 1:189–200, 1992.
- [AES10] Boris Aronov, Esther Ezra, and Micha Sharir. Small-size ϵ -nets for axis-parallel rectangles and boxes. *SIAM J. Comput.*, 39(7):3248–3282, 2010.
- [AES12] Pankaj K. Agarwal, Esther Ezra, and Micha Sharir. Near-linear approximation algorithms for geometric hitting sets. *Algorithmica*, 63(1-2):1–25, 2012.
- [AGK⁺98] Sanjeev Arora, Michelangelo Grigni, David R. Karger, Philip N. Klein, and Andrzej Woloszyn. A polynomial-time approximation scheme for weighted planar graph tsp. In *SODA*, pages 33–41, 1998.
- [AGK⁺01] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristic for k-median and facility location problems. In *STOC*, pages 21–29, 2001.
- [AK92] N. Alon and D. Kleitman. Piercing convex sets and the Hadwiger Debrunner (p; q)-problem. *Adv. Math.*, 96(1):103–112, 1992.
- [Alo12] Noga Alon. A non-linear lower bound for planar epsilon-nets. *Discrete & Computational Geometry*, 47(2):235–244, 2012.
- [ALST03] G. Aloupis, S. Langerman, M. A. Soss, and G. T. Toussaint. Algorithms for bivariate medians and a Fermat-Torricelli problem for lines. *Comput. Geom.*, 26(1):69–79, 2003.
- [Ame94] Nina Amenta. Helly-type theorems and generalized linear programming. *Discrete & Computational Geometry*, 12:241–261, 1994.
- [AvKS98] Pankaj K. Agarwal, Mark van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. *Computational Geometry: Theory and Applications*, 11:209–218, 1998.
- [AW13] A. Adamaszek and A. Wiese. Approximation schemes for maximum weight independent set of rectangles. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*, 2013.
- [AW14] A. Adamaszek and A. Wiese. A QPTAS for maximum weight independent set of polygons with polylogarithmically many vertices. In *SODA*, 2014.
- [Bár82] Imre Bárány. A generalization of Carathéodory’s theorem. *Discrete Mathematics*, 40(2-3):141–152, 1982.
- [BDMR01] Piotr Berman, Bhaskar DasGupta, S. Muthukrishnan, and Suneeta Ramaswami. Efficient approximation algorithms for tiling and packing problems with rectangles. *Journal of Algorithms*, 41:443–470, 2001.

- [BF82] E Boros and Z Furedi. The number of triangles covering the center of an n -set. *Geometriae Dedicata*, 17(1):69–77, 1982.
- [BG95] H. Bronnimann and M. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.
- [Bla23] W. Blaschke. *Vorlesungen uber Differentialgeometrie. II, Affine Differentialgeometrie*. Springer, 1923.
- [BMN09] Boris Bukh, Jirí Matousek, and Gabriel Nivasch. Lower bounds for weak epsilon-nets and stair-convexity. In *Symposium on Computational Geometry*, pages 1–10, 2009.
- [BMN10] Boris Bukh, Jirí Matousek, and Gabriel Nivasch. Stabbing simplices by points and flats. *Discrete & Computational Geometry*, 43(2):321–338, 2010.
- [Buk06] B. Bukh. A point in many triangles. *Electronic Journal of Combinatorics*, 13(10), 2006.
- [CC09] Parinya Chalermsook and Julia Chuzhoy. Maximum independent set of rectangles. In *SODA*, pages 892–901, 2009.
- [CDD⁺09] Francisco Claude, Reza Dorrigiv, Stephane Durocher, Robert Fraser, Alejandro López-Ortiz, and Alejandro Salinger. Practical discrete unit disk cover using an exact line-separable algorithm. In *ISAAC*, pages 45–54, 2009.
- [CDI⁺13] Dan Chen, Olivier Devillers, John Iacono, Stefan Langerman, and Pat Morin. Oja centers and centers of gravity. *Comput. Geom.*, 46(2):140–147, 2013.
- [CEG⁺93] Bernard Chazelle, Herbert Edelsbrunner, Michelangelo Grigni, Leonidas Guibas, Micha Sharir, and Emo Welzl. Improved bounds on weak ε -nets for convex sets. In *Proceedings of the twenty-fifth annual ACM Symposium on Theory of Computing*, pages 495–504, 1993.
- [CF90] Bernard Chazelle and Joel Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- [CF13] Victor Chepoi and Stefan Felsner. Approximating hitting sets of axis-parallel rectangles intersecting a monotone curve. *Comput. Geom.*, 46(9):1036–1041, 2013.
- [CGKS12] Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *SODA*, 2012.
- [Cha03] Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46:178–189, 2003.
- [CHP09] Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. In *Symposium on Computational Geometry*, pages 333–340, 2009.
- [CKLT07] P. Carmi, M. Katz, and N. Lev-Tov. Covering points by unit disks of fixed location. In *ISAAC*, pages 644–655, 2007.
- [CMWZ04] Gruia Calinescu, Ion I. Mandoiu, Peng-Jun Wan, and Alexander Z. Zelikovsky. Selecting forwarding neighbors in wireless ad hoc networks. *Mob. Netw. Appl.*, 9(2):101–111, 2004.
- [CV07] K. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete Comput. Geom.*, 37:43–58, 2007.

- [DFLON11] Gautam K. Das, Robert Fraser, Alejandro López-Ortiz, and Bradford G. Nickerson. On the discrete unit disk cover problem. In *WALCOM*, pages 146–157, 2011.
- [Die00] Reinhard Diestel. *Graph theory*. Springer-Verlag, New York, 2000.
- [EJS01] Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric graphs. In *Proc. of the 12th Annual ACM-SIAM Sympos. on Discrete Algorithms*, pages 671–679, 2001.
- [ERS05] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the VC-dimension is small. *Inf. Process. Lett.*, 95:358–362, 2005.
- [EvL08] Thomas Erlebach and Erik Jan van Leeuwen. Approximating geometric coverage problems. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, 2008.
- [FGL⁺11] J. Fox, M. Gromov, V. Lafforgue, A. Naor, and J. Pach. Overlap properties of geometric expanders. In *SODA*, 2011.
- [FHP08] R. Fulek, A. Holmsen, and J. Pach. Intersecting convex sets by rays. In *Proceedings of the Symposium on Computational Geometry*, pages 385–391, 2008.
- [FP09] J. Fox and J. Pach. Separator theorems and turan-type results for planar intersection graphs. *Advances in Mathematics*, pages 1070–1080, 2009.
- [FP11] Jacob Fox and János Pach. Computing the independence number of intersection graphs. In *SODA*, pages 1161–1165, 2011.
- [Fre87] G. N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.
- [GKKV09] Matt Gibson, Gaurav Kanade, Erik Krohn, and Kasturi R. Varadarajan. An approximation scheme for terrain guarding. In *APPROX-RANDOM*, pages 140–148, 2009.
- [GLS81] M. Grotschel, L. Lovász, and A. Schrijver. Polynomial algorithms for perfect graphs. *Ann. Discrete Math.*, 21:211–356, 1981.
- [Gol04] Martin Charles Golumbic. *Algorithmic Graph Theory and Perfect Graphs (Annals of Discrete Mathematics, Vol 57)*. North-Holland, 2004.
- [GP10] Matt Gibson and Imran A. Pirwani. Approximation algorithms for dominating set in disk graphs. *CoRR*, abs/1004.3320, 2010.
- [Gro10] M. Gromov. Singularities, expanders and topology of map. part 2: From combinatorics to topology via algebraic isoperimetry. *Geom. Func. Anal.*, 20:416–526, 2010.
- [HM84] Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in robotics and vlsi. In *STACS*, pages 55–62, 1984.
- [HM87] D. S. Hochbaum and W. Maass. Fast approximation algorithms for a nonconvex covering problem. *J. Algorithms*, 8(3):305–323, 1987.
- [HMR⁺98] Harry B. Hunt, Madhav V. Marathe, Venkatesh Radhakrishnan, S. S. Ravi, Daniel J. Rosenkrantz, and Richard Edwin Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26:238–274, 1998.

- [HW87] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [IA83] H. Imai and Ta. Asano. Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. *J. Algorithms*, 4:310–323, 1983.
- [Kar12] R.N. Karasev. A simpler proof of the boros-furedi-barany-pach-gromov theorem. *Discrete and Computational Geometry*, 47(3):492–495, 2012.
- [KMN⁺02] Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for k-means clustering. In *Symposium on Computational Geometry*, pages 10–18, 2002.
- [KMS12] D. Kral, L. Mach, and J. S. Sereni. A new lower bound based on Gromov’s method of selecting heavily covered points. *Discrete and Computational Geometry*, to appear, 2012.
- [KPW92] J. Komos, J. Pach, and G. Woeginger. Almost tight bounds for epsilon nets. *Discrete & Computational Geometry*, pages 163–173, 1992.
- [LENO02] Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. Routing and admission control in networks with advance reservations. In *Intl. Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 215–228, 2002.
- [Lov74] L. Lovász. Problem 206. *Matematikai Lapok*, 25, 1974.
- [LT77] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. Technical report, Stanford University, Stanford, CA, USA, 1977.
- [LT80] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [Mat92] Jirí Matousek. Efficient partition trees. *Discrete & Computational Geometry*, 8:315–334, 1992.
- [Mat02] Jirí Matousek. *Lectures in Discrete Geometry*. Springer-Verlag, New York, NY, 2002.
- [Mil86] Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.*, 32(3):265–279, 1986.
- [MSW90] J. Matousek, R. Seidel, and E. Welzl. How to net a lot with little: Small epsilon-nets for disks and halfspaces. In *Proceedings of Symposium on Computational Geometry*, pages 16–22, 1990.
- [MW04] Jirí Matousek and Uli Wagner. New constructions of weak epsilon-nets. *Discrete & Computational Geometry*, 32(2):195–206, 2004.
- [MW10] J. Matousek and U. Wagner. On gromov’s method of selecting heavily covered points. In *Manuscript*, 2010.
- [NV06] S. Narayanappa and P. Vojtechovský. An improved approximation factor for the unit disk covering problem. In *CCCG*, 2006.
- [Oja83] H. Oja. Descriptive statistics for multivariate distributions. *Statistics and Probability Letters*, 1:327–332, 1983.

- [PA95] Janos Pach and Pankaj K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, New York, NY, 1995.
- [PR08] E. Pyrga and S. Ray. New existence proofs for epsilon-nets. In *Proceedings of Symposium on Computational Geometry*, pages 199–207, 2008.
- [PT11] János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Symposium on Computational Geometry*, pages 458–463, 2011.
- [PW90] J. Pach and G. Woeginger. Some new bounds for epsilon-nets. In *Symposium on Computational Geometry*, pages 10–15, 1990.
- [Rad47] R. Rado. A theorem on general measure. *J. London. Math. Soc.*, 21:291–300, 1947.
- [RN95] C. S. Rim and K. Nakajima. On rectangle intersection and overlap graphs. *IEEE Trans. on Circuits and Systems*, 42(9):549–553, 1995.
- [ROO03] T. Ronkainen, H. Oja, and P. Orponen. Computation of the multivariate oja median. In *R. Dutta and P. Filzmoser. International Conference on Robust Statistics*, 2003.
- [RR96] P. Rousseeuw and I. Ruts. Bivariate location depth. *Applied Statistics*, 45:516–526, 1996.
- [SSW08] Shakhbar Smorodinsky, Marek Sulovský, and Uli Wagner. On center regions and balls containing many points. In *Proceedings of the 14th annual International Conference on Computing and Combinatorics, COCOON '08*, pages 363–373, 2008.
- [Tuk75] J. Tukey. Mathematics and the picturing of data. In *Proc. of the international congress of mathematicians*, pages 523–531, 1975.
- [TV93] H. Tverberg and S. Vrecica. On generalizations of Radon’s theorem and the Ham Sandwich theorem. *Europ. J. of Combinatorics*, 14:259–264, 1993.
- [Var10] Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the 42nd ACM symposium on Theory of computing, STOC '10*, pages 641–648, 2010.
- [Wag03] U. Wagner. *On k -Sets and Applications*. PhD thesis, ETH Zurich, 2003.