Weak and Strong ϵ-Nets for Geometric Range Spaces

Saurabh Ray

Dissertation

Presented to the Faculties

of Natural Science and Technology

of Saarland University

in Candidacy for the degree of

Doctor of Engineering Sciences (Dr.-Ing)

Dissertation

zur Erlangung des Grades

des Doktors der Ingenieurswissenschaften (Dr.-Ing)

der Naturwissenschaftlich-Technischen Fakultäten

der Universität des Saarlandes

Saarbrücken

February 2009

Defense Details: Saurabh Ray

Date of Defense	:	25th May, 2009
Defense Chair	:	Prof. Dr. Jens Dittrich
Defense Committee	:	Prof. Dr. Raimund Seidel
	:	Prof. Dr. Kurt Mehlhorn
	:	Prof. Dr. Janos Pach
Minutes Writer	:	Rajiv Raman, PhD

I hereby declare that both the electronic and the printed versions of this thesis that have been submitted for obtaining the Doctorate certicate, match exactly in terms of content.

Saarbrücken, May 25th 2009

Saurabh Ray

Betreunder Hochschullehrer - Supervisor

Prof. Dr. Raimund Seidel,

Universität des Saarlandes, Saarbrücken, Germany

Gutachter - Reviewers

Prof. Dr.h.c. Kurt Mehlhorn,

Max-Plank-Institut für Informatik, Saarbrücken, Germany

Prof. Janos Pach,

Ecole Polytechnique Fédérale de Lausanne, Switzerland

Dekan - Dean

Prof. Dr. Joachim Weickert,

Universität des Saarlandes, Saarbrücken, Germany

Datum des Kolloquiums - Date of Defense

Saurabh Ray,

FR 6.2 Informatik,

Universität des Saarlandes,

Saarbrücken, Germany

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, dass ich die vorliegende Arbeit selbständig und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe. Die aus anderen Quellen oder indirekt ubernommenen Daten und Konzepte sind unter Angabe der Quelle gekennzeichnet.

Die Arbeit wurde bisher weder im In- noch im Ausland in gleicher oder ähnlicher Form in einem Verfahren zur Erlangung eines akademischen Grades vorgelegt.

Saarbrücken, den 19.02.2009

(Unterschrift)

Kurzfassung

Diese Arbeit beschäftigt sich mit ϵ-nets in der Geometrie und verwandten Problemen. Im ersten Teil der Arbeit werden starke ϵ-nets und das eng verwandte Minimum Hitting Set Problem betrachtet. Es wird eine neue Technik vorgestellt mit deren Hilfe die Existenz von kleinen ϵ-nets in verschiedenen geometrischen Bereichsräumen nachgewiesen werden kann. Diese Technik liefert auch effiziente Algorithmen um kleine ϵ-nets zu berechnen. Mit der bekannten Reduktion von Bronimann und Goodrich [10], führt dies zu Approximationsalgorithmen mit konstantem Faktor für die entsprechenden Hitting Set Probleme. Der Approximationsfaktor kann sogar verbessert werden durch einen relative einfachen, auf lokaler Suche basierenden Ansatz, der zu dem ersten polynomiellen Approximationsschema führt.

Der zweite Teil der Arbeit ist den schwachen ϵ-nets gewidmet die eine wichtige Verallgemeinerung der starken ϵ-nets in konvexen Bereichen darstellen. Zunächst wird der einfachste Fall der schwachen ϵ-nets betrachtet, der Centerpoint. Es wird ein neuer, einfacherer Beweis für das bekannte Centerpoint Theorem (und ebenso Helly's Theorem) in beliebiger Dimension gezeigt. Die gleiche Idee lässt sich auch benutzen um eine optimale Verallgemeinerung der Centerpoints zu zwei Punkten in der Ebene zu zeigen. Mit dieser Technik können verschiedene Resultate für schwache ϵ-nets in der Ebene verbessert werden. Abschlieend wird das allgemeine schwache ϵ-net Problem in d Dimensionen betrachtet. Eine langjährige Vermutung besagt, dass schwache ϵ-nets der Grösse $O(\epsilon^{-1}\text{polylog}\epsilon^{-1})$ für konvexe Mengen in jeder Dimension existieren. Es stellt sich heraus, dass wenn sich die Vermutung als wahr erweist, dann ist es möglich ein schwaches ϵ-net aus einer kleinen Menge von Inputpunkten zu erzeugen. In dieser Arbeit wird gezeigt, dass dies tatsächlich möglich ist und ein schwaches ϵ-net aus $O(\epsilon^{-1}\mathrm{polylog}\epsilon^{-1})$ Inputpunkten erzeugt werden kann. Letztendlich lässt sich ein interessanter Zusammenhang zwischen schwachen und starken ϵ-nets zeigen durch den schwache ϵ-nets durch eine Zufallsauswahl konstruiert werden können.

Abstract

This thesis deals with strong and weak ϵ-nets in geometry and related problems. In the first half of the thesis we look at strong ϵ-nets and the closely related problem of finding minimum hitting sets. We give a new technique for proving the existence of small ϵ-nets for several geometric range spaces. Our technique also gives efficient algorithms to compute small ϵ-nets. By a well known reduction due to Bronimann and Goodrich [10], our results imply constant factor approximation algorithms for the corresponding minimum hitting set problems. We show how the approximation factor given by this standard technique can be improved by giving the first polynomial time approximation scheme for some of the minimum hitting set problems. The algorithm is a very simple and is based on local search.

In the second half of the thesis, we turn to weak ϵ-nets, a very important generalization of the idea of strong ϵ-nets for convex ranges. We first consider the simplest example of a weak ϵ-net, namely the centerpoint. We give a new and arguably simpler proof of the well known centerpoint theorem (and also Helly's theorem) in any dimension and use the same idea to prove an optimal generalization of the centerpoint to two points in the plane. Our technique also gives several improved results for small weak ϵ-nets in the plane. We finally look at the general weak ϵ-net problem is d-dimensions. A long standing conjecture states that weak ϵ-nets of size $O(\epsilon^{-1}\text{polylog}\epsilon^{-1})$ exist for convex sets in any dimension. It turns out that if the conjecture is true then it should be possible to construct a weak ϵ-net from a small number of input points. We show that this is indeed true and it is possible to construct a weak ϵ-net from $O(\epsilon^{-1}\text{polylog}\epsilon^{-1})$ input points. We also show an interesting connection between weak and strong ϵ-nets which shows how random sampling can be used to construct weak ϵ-nets.

Acknowledgements

Many people have directly or indirectly contributed to this thesis. I would like to thank my supervisor Prof. Raimund Seidel for his guidance and support. He was always available to listen to my problems and ideas and showed great confidence in me. Thanks also to his great database of counterexamples, a lot of effort which would have gone in fruitless directions was saved. I would like to thank Nabil Mustafa for not only introducing me to the exciting area of discrete geometry but also for teaching me how to think and solve problems and how to enjoy research. I have not learnt from any other source as much as I have learnt from working with Nabil. I wish to thank my coauthors and friends Hansraj Tiwary, Deepak Ajwani, Evangelia Pyrga, Khaled Elbassioni, Sathish Govindarajan and Rajiv Raman for the countless exciting brainstorming sessions from which I have learnt many technical as well as non-technical things. Thanks to MPII, Saarbrücken, I had access to so many researchers. Hansraj was also a great office mate. He always showed a lot of excitement in any non-trivial topic and provided me with many interesting things to think about. Without him, the time spent in the office would have been very dry. I am also very lucky to have friends like Imran Rauf, Ralitsa Angelova, Ibne Hassan, Abdul Qadir Kara, Georgiana Ifrim, Sajjad Hussain and Waqar Saleem who made life in Saarbrücken very enjoyable. I am grateful to Prof. Kurt Mehlhorn and Prof. Janos Pach for agreeing to examine my thesis. Finally, I would like to thank my parents, my sister and my wife who make my life meaningful.

Contents

List of Figures

Chapter 1

Introduction

This thesis deals with the strong and weak ϵ-nets, which are fundamental tools in discrete and computational geometry. The theory of strong ϵ-nets has been very successful and has found many applications in computational geometry, statistics and learning theory. In computational geometry, strong ϵ-nets (and ϵ-approximations) power many data structures and algorithms used in point location, range searching, range counting and several other tools for geometric divide and conquer. They also find use in derandomizing divide and conquer type algorithms. The idea of strong ϵ-nets was extended to weak ϵ-nets for convex sets by Haussler and Welzl in their seminal paper [26]. Weak ϵ-nets have found application, among other things, in the beautiful proof of the Hadwiger-Debrunner (p,q) conjecture by Alon and Kleitman [6]. In this chapter, we give a brief introduction to the idea of strong and weak ϵ-nets and present some of the most important results in this area. We then give a summary of the work presented in this thesis.

1.1 Strong ϵ-nets

A range space is a set system $\mathcal{R} = (X, R)$ where X is a (possibly infinite) set called the *ground set* and R is the set of subsets of X. We will call the elements of X the *points* and the elements of R the *ranges* of $\mathcal{R}$. In this thesis we will deal only with range spaces with finite ground sets. Given a parameter $0 < \epsilon < 1$, we say that a range $r \in R$ is *ϵ-heavy* if $|r| > \epsilon|X|$. A *strong* ϵ-net for $\mathcal{R}$ is a subset $Y \subseteq X$ which *hits* all ϵ-heavy ranges i.e. it has a non-empty intersection with each ϵ-heavy range in R. In the following we will just write "ϵ-net" for "strong ϵ-net". We are interested in small ϵ-nets since in some sense they allows us to approximate the given range space economically. Since each of the ranges we want to hit are ϵ-heavy, if we randomly pick a point from $\mathcal{R}$, we hit any given ϵ-heavy range with probability at least ϵ. Since there are at most $|R|$ ranges to be hit, a simple calculation shows that a random sample of X of size $\frac{1}{\epsilon}\log|R|$ hits all the ranges simultaneously with positive probability. Hence we have:

Proposition 1. *Any finite range space $\mathcal{R} = (X, R)$ admits an ϵ-net of size $O(\frac{1}{\epsilon}\log|R|)$.*

It turns out that in many cases it is possible to get a better upper bound. For example, consider the range space $\mathcal{R} = (X, R)$ in which the ground set X is a finite set of n points in the plane and the ranges in R are the subsets of X which can be obtained by intersecting X with some triangle in the plane. In this case, the number of ranges is at most $O(n^6)$ since we can always change a triangle, without changing the set of points it contains, so that each of its sides passes through two of the points. Proposition 1 therefore guarantees an ϵ-net of size $\frac{1}{\epsilon}\log n$. However, it can be shown that there is an ϵ-net of size $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ for this range space. Since ϵ is typically a constant, $\frac{1}{\epsilon}$ can be assumed to be smaller than n. This, therefore, gives a bound better than Proposition 1.

Such small ϵ-nets obviously do not exist for all range spaces. For example, if the ground set X is a set of points in the plane lying in convex position and the ranges are subsets of X that can be obtained by intersecting X with a convex set. In this case, the ranges consists of all subsets of X and hence any ϵ-net has to have $|X| - (\epsilon|X| + 1)$ points since if we leave out any set of $\epsilon|X| + 1$ points then those points together form an ϵ-heavy range that is not hit.

What leads to an improved upper bound in the case of triangular ranges is the fact that not only does $\mathcal{R}$ have a polynomial number of ranges, the number of ranges is *hereditarily* polynomial i.e. the number of distinct ranges induced by any subset of $Z \subseteq X$, $|\{r \cap Z | r \in R\}|$ is polynomial in $|Z|$. One way to capture this property is the notion of the *Vapnik Chervonenkis* dimension (VC dimension in short) of a range space which was introduced by Vapnik and Chervonenkis in [49]. Given a range space $\mathcal{R} = (X, R)$, the *projection* of $\mathcal{R}$ on a subset $Y \subseteq X$ is defined as $\mathcal{R}|_Y = \{Y \cap r | r \in R\}$. We say that a set A is shattered by $\mathcal{R}$ if all subsets of A can be obtained by intersecting A with some range in R i.e. $\mathcal{R}|_A = 2^A$. The VC dimension of $\mathcal{R}$, denoted $dim(\mathcal{R})$, is the size of the largest set shattered by $\mathcal{R}$. It has be shown that range spaces of finite VC dimension are precisely the ones with the hereditarily polynomial property [49, 47].

Haussler and Welzl [26], who introduced the notion of ϵ-nets, showed that range spaces with a small VC dimension admit a small ϵ-net. More precisely, they show that

Theorem 1 (ϵ-net theorem)**.** *For any finite range space $\mathcal{R} = (X, R)$ with $dim(\mathcal{R}) \leq d$ and parameters $0 < \epsilon, \delta < 1$, a random subset $N \subseteq X$ of size $\max\left\{\frac{8d}{\epsilon} \log \frac{8d}{\epsilon}, \frac{4}{\epsilon} \log \frac{2}{\delta}\right\}$ is an ϵ-net for $\mathcal{R}$ with probability at least $1 - \delta$.*

Komlós et al. [31] have also shown that this is tight up to constant factors. The constant factor was improved by Blumer et al. [9]. Komlós et al. [32] improved it further and show the following:

Theorem 2. *Let $f(d, \epsilon)$ denote the maximum size, over all range spaces of VC dimension at most d, of an ϵ-net of the smallest size for that range space. Then, $f(d, \epsilon) = (1+o(1))\frac{d}{\epsilon} \log \frac{1}{\epsilon}$.*

Vapnik and Chervonenkis [49] also introduced the notion of *ϵ-approximations*. The purpose of an ϵ-net is to hit all ϵ-heavy ranges. The purpose of an ϵ-approximation is something stronger. We want the fraction of the points of the ϵ-approximation which lie inside each range to be equal to the relative size of the range within an additive factor of ϵ. More precisely, given a range space $\mathcal{R} = (X, R)$, $A \subseteq X$ is an ϵ-approximation for $\mathcal{R}$ if for every range $r \in R$,

$$\left| \frac{|A \cap r|}{|A|} - \frac{|X \cap r|}{|X|} \right| \leq \epsilon.$$

As for ϵ-nets, they show that range space with a finite VC dimension d also has a small ϵ-approximation [49]:

Theorem 3. *Any range space of VC dimension d admits and ϵ-approximation of size $O(d/\epsilon^2 \log (d/\epsilon))$.*

Matousek et al. [38] improved the above bound to $O(\epsilon^{-(2-2/(d+1))} \log^{2-2/(d+1)} \epsilon^{-1})$ and showed that this is almost tight up to polylogarithmic factors.

A notion closely related to the notion of VC dimension is that of *scaffold dimension* [24] which is based on the shatter function $\pi_{\mathcal{R}}(m)$ of a range space $\mathcal{R} = (X, R)$ that denotes the maximum number of distinct ranges in $\mathcal{R}|_Y$ for an m-sized subset Y of X i.e. $\pi_{\mathcal{R}}(m) = max \{ \ |\mathcal{R}|_Y| \ : \ Y \subseteq X, \ |Y| = m \ \}$. The scaffold dimension of $\mathcal{R}$ is the smalled integer d such that $\pi_{\mathcal{R}}(m) = O(m^d)$. It has been shown that the scaffold dimension of a range space is always at most its VC dimension [49, 47, 24]. Also, if the scaffold dimension of a range space is d, then clearly a very large set cannot be shattered. The number of distinct ranges induced by a shattered set of size m is 2^m and for large enough m this is $\omega(m^d)$. Therefore, a finite scaffold dimension implies a finite VC dimension. The scaffold dimension is often easier to work with since for many ranges spaces, especially those arising in geometry, it is easier to prove that a range space has a bounded scaffold dimension.

Matousek [39] gave efficient deterministic algorithms for computing ϵ-approximations and ϵ-nets for range spaces $\mathcal{R} = (X, R)$ with a finite scaffold dimension. The algorithms assume the existence of a *subspace oracle* i.e. an oracle which given a set Y enumerates the ranges in $\mathcal{R}|_Y$ in time $O(|Y|\, k)$ where k is the number of ranges enumerated. He showed that

Theorem 4. *Given a subspace oracle for a range space $\mathcal{R} = (X, R)$ of scaffold dimension $d > 1$ and a parameter $0 < \epsilon < 1$, we can deterministically compute an ϵ-approximation of size $O(\epsilon^{-2} \log \epsilon^{-1})$ and an ϵ-net of size $O(\epsilon^{-1} \log \epsilon^{-1})$ in time $O(|X| \cdot (\epsilon^{-2} \log \epsilon^{-1})^d)$ time.*

Matousek [39] also proved the same result for weighted sets, i.e. when the ground set X is equipped with a probability measure and the weight of a range is the total weight of the elements in it.

In this thesis, we will be mostly concerned with range spaces induced by a set of points and a set of geometric objects. We have already encountered such a range space, namely the one in which the ground set is a finite set P of points in the plane and the ranges are defined by the set of all triangles in the plane. The range defined by each triangle is the set of points contained in it. It follows from the arguments before that this range space has a finite scaffold dimension and hence admits and ϵ-net of size $O(\epsilon^{-1} \log \epsilon^{-1})$. One natural question for such geometric range spaces is whether we can get a smaller ϵ-net by exploiting the geometric structure. This is the question that we address in Chapter 2. We consider two main kinds of range spaces induced by a set P of points and a set S of geometric objects. In the first kind, we treat the points as the ground set, and each geometric object $s \in S$ defines the range $s \cap P$. We call this the *primal range space induced by P and S*. In the second kind, we exchange the roles of P and S i.e., S is the ground set and each $p \in P$ defines the range $\{s \in S : p \in s\}$. Several people have proved the existence of ϵ-nets of size $O(1/\epsilon)$ for such geometric range spaces, improving on the $O(\epsilon^{-1} \log \epsilon^{-1})$ bound for range spaces of finite VC dimension. Pach and Woeginger [44] proved that halfspaces and translates of polytopes in $\mathbb{R}^2$ admit strong ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek et al. [36] proved that halfspaces in $\mathbb{R}^3$ and certain special families of pseudo-disks in $\mathbb{R}^2$ (they require that there is exactly one pseudo-disk through any three non-collinear points in the plane) admit strong ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek later found a shorter proof for the existence of $O(\frac{1}{\epsilon})$ size strong ϵ-nets for halfspaces in $\mathbb{R}^3$ via shallow cuttings [34]. Clarkson and Varadarajan [16] gave a fairly general framework which applies to many such geometric range spaces. Their technique is particularly useful for the dual range spaces where they show that if the geometric objects under consideration have a *small* union complexity, then is possible to get a correspondingly small ϵ-net. Their technique however does not readily apply to primal range spaces.

In Chapter 2, we develop a new technique for proving the existence of $o(\epsilon^{-1} \log \epsilon^{-1})$ size ϵ-nets and use it to obtain a $O(1/\epsilon)$ size ϵ-net for the primal range space induced by a set of points and a set of pseudo-disks in the plane. This result was not previously known and it is not clear whether the technique of Clarkson and Varadarajan [16] can be used for this case. Our technique also gives a new proof for almost all geometric range spaces for which such a small ϵ-net is known to exist. In particular, it gives a very short and elementary proof, using only double counting arguments, for the existence of $O(1/\epsilon)$ for the range space induced by a set of points and halfspaces in $\mathbb{R}^3$. The earlier proofs used fairly sophisticated geometric and/or probabilistic tools. The proof technique used also implies fast algorithms for computing ϵ-nets of small size. In the process of applying our technique to the various geometric range spaces, we also prove interesting combinatorial results about them. For example, we prove that given a set of points and a set of pseudo-disks in the plane, it is possible to construct a planar graph whose vertices are the set of points so that the subgraph induced by the vertices in any of the given pseudo-disks is connected. Such results lead to a PTAS for the hitting sets problems related to these range spaces. We describe the hitting set problem and the specific results obtained in next section (Section 1.2).

1.2 Hitting Sets

The problem of computing *minimum hitting sets* is very closely related to the problem of computing small ϵ-nets. In the (strong) ϵ-net problem we are interested in finding a small subset of the ground set which *hits*, i.e. has a non-empty intersection with, all the ϵ-heavy ranges. A natural algorithmic question then is whether we can compute, exactly or approximately, an ϵ-net of the smallest size. A more general question is if we are given a set of ranges (which may or may not be ϵ-heavy), can we compute the smallest *hitting set* i.e., the smallest subset of the ground set which has a non-empty intersection with each of the given ranges? If the ranges are allowed to be arbitrary subsets of the ground set, then this problem is the same as the set cover problem whose approximability is completely resolved. It is possible to get a $O(\log n)$ approximation, where n is the size of the ground set, using a greedy algorithm and it is not possible to do better unless P=NP. However, when the range space is "simple", it is often possible to get better approximations. Bronimann and Goodrich [10] reduced the problem of computing hitting sets to computing weighted ϵ-nets i.e. hitting ϵ-heavy ranges when the elements of the ground set have non-uniform weights. Essentially, they compute weights so that each of given ranges are δ-heavy for as large δ as possible and then compute a δ-net. They show that since one can compute $O(\epsilon^{-1} \log \epsilon^{-1})$ size ϵ-nets for range spaces of finite VC dimension one gets hitting sets of size $O(\textsc{opt} \log \textsc{opt})$, where $\textsc{opt}$ is the size of the smallest hitting set, via this reduction. In many cases, as we will see in Chapter 2, it is possible to compute an ϵ-net of size $O(1/\epsilon)$. In such cases, the reduction gives a constant factor approximation to the smallest hitting set. If we can compute an ϵ-net of size c/ϵ, then we get a c-approximation. The weakness of this approach is that after reducing the minimum hitting set problem to a problem of computing an ϵ-net, we use a worst case bound for the ϵ-net. However, the worst case bound can be much worse than the smallest ϵ-net for the problem at hand. Hence, c can be quite large. Even for simple ranges like halfspaces in the plane, c is at least 2 [43] and this rules out the possibility of a PTAS using this approach. In Chapter 3, we give a new general technique for approximating geometric hitting sets, which avoids the limitation of the Bronnimann-Goodrich technique. We give the first polynomial-time approximation schemes for the minimum geometric hitting set problem for a wide class of geometric range spaces. All these problems are strongly NP-complete and hence, unless P=NP, there is no FPTAS for these problem. Specifically, we show that:

- Given a set P of n points, and a set $\mathcal{H}$ of m halfspaces in $\mathbb{R}^3$, one can compute a $(1+\delta)$-approximation to the smallest subset of P that hits all the halfspaces in $\mathcal{H}$ in

$O(mn^{O(\delta^{-2})})$ 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 $(1+\delta)$-approximation to the smallest subset of P that hits all the regions in $\mathcal{D}$ in $O(mn^{O(\delta^{-2})})$ time. This includes pseudo-disks (they are 2-admissible), same-height rectangles, circular disks, translates of convex objects etc. See Definition 3.3.1 for the definition of an r-admissible set of regions.

The above results should be contrasted with the fact that even for relatively simple range spaces like those induced by unit disks in the plane, the previous best known approximation algorithm is due to a recent paper of Carmi *et. al.* [11] which gives a 38-approximation algorithm improving the earlier best known factor of 72 [42].

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 set $k = c\delta^{-2}$, where c is a constant, to get a $(1+\delta)$-approximation algorithm.

In order to prove that the above local search algorithm works we use some combinatorial results about these geometric ranges that we derive in Chapter 2 in conjunction with a theorem about planar bipartite graphs (Theorem 18) which we consider to be interesting in its own right. It states the following:

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.

As a side effect of the above theorem, we also get a PTAS for computing the maximum independent sets of the intersection graph of a given set of r-admissible regions. This new result extends the results obtained in [4] and [1].

1.3 Weak ϵ-Nets

A natural question for the primal range spaces induced by a set of points and set of geometric objects is whether it is possible to get a smaller ϵ-net if we do not insist that the ϵ-net be a subset of the point set P i.e. if we allow it to be an arbitrary subset of the space in which they are embedded (e.g. $\mathbb{R}^d$). Such subsets are called *weak* ϵ-nets in order to distinuish them from the strong ϵ-nets we have considered so far. In some cases one can indeed get a very small weak ϵ-net. This for example is the case when the ranges are defined by half-spaces in the plane. The three corners of any large enough triangle containing the set P is a weak ϵ-net. On the other hand, even for simple geometric objects like triangles, it is not known whether one can do any better than finding a strong ϵ-net of size $O(\epsilon^{-1} \log \epsilon^{-1})$. Perhaps the most interesting case is the range space induced by convex sets which does not have a finite VC dimension and does not admit a small strong ϵ-net. It turns out that small weak ϵ-nets do exist for this range space!

The concept of weak ϵ-nets with respect to convex ranges was introduced by Haussler and Welzl [26] and the notion has found several applications in discrete and combinatorial geometry (see Matousek's book for several examples [35]).

Let $w(d, \epsilon)$ denote the maximum size of the weak ϵ-net required for any set of points in $\mathbb{R}^d$ under convex ranges. Alon *et al.* [5] have shown that it is finite and for any ϵ, d, there exist a weak ϵ-net whose size is independent of size of the ground set. Specifically, they proved that $w(d, \epsilon) \leq O(1/\epsilon^{d+1-\delta_d})$, where δ_d tends to zero with $d \to \infty$. They also showed that for a set of points in $\mathbb{R}^2$ in convex position, there exists a weak ϵ-net of size $O(1/\epsilon \operatorname{polylog}(1/\epsilon))$. More recently, Matousek and Wagner [37] gave an elegant algorithm that computes weak ϵ-nets in $\mathbb{R}^d$ of size $O(1/\epsilon^d \operatorname{polylog}(1/\epsilon))$.

One special case of weak ϵ-nets is the *centerpoint*. When ϵ is large enough the weak ϵ-net consists of just one point which is called the centerpoint. This special case is well studied and the famous centerpoint theorem [43, 35] states that if we have points in $\mathbb{R}^d$ and $\epsilon > d/(d+1)$, all ϵ-heavy convex ranges can be *hit* by just one point.

While the situation with strong ϵ-nets is very well understood, our understanding of weak ϵ-nets is far from satisfactory. The best upper bound known for the size of weak ϵ-nets in $\mathbb{R}^d$ is $O(1/\epsilon^d)$ although it is not clear why it should be significantly larger than $\frac{1}{\epsilon}\operatorname{polylog}\frac{1}{\epsilon}$. Matousek and Wagner [37] have conjectured that $O(\frac{1}{\epsilon}\operatorname{polylog}\frac{1}{\epsilon})$ is the right upper bound. This remains one of the most important open problems in the area.

In Chapter 4, we study weak ϵ-nets of a small constant size. We start by looking at an alternate proof of the centerpoint theorem. The centerpoint theorem is usually proved by using Helly's theorem which in turn is proved by using Radon's theorem. We give a very short proof of the centerpoint theorem using an elementary argument which avoids using Helly's theorem and Radon's theorem. The same idea gives a simple proof of Helly's theorem too. We then prove that in the plane, given n points, it is possible to pick two points p and q in the plane (not necessarily among input points) so that any convex set containing more than $4n/7$ input points contains at least one of the two points p and q. We also show that this is tight i.e. it is not possible to pick two points which *hit* all convex sets containing at least $4n/7$ points. This gives an optimal extension of the centerpoint theorem to two points in the plane. We finally look at the cases of 3 or more points and improve several bounds obtained by Aronov *et al.* [8].

In Chapter 5 we turn to the general weak ϵ-net problem with respect to convex ranges in $\mathbb{R}^d$. As we have remarked before the bounds for the size of weak ϵ-nets are not very satisfactory and there hasn't been any progress in a long time. We consider the conjecture of Matousek and Wagner [37] which states that $O(\frac{1}{\epsilon}\text{polylog}\frac{1}{\epsilon})$ is the right upper bound. We then make the following observation (Observation): Given a set P of n points in $\mathbb{R}^d$, a weak ϵ-net of P of size k is completely described by $O(d^2k)$ points of P. For example in the plane, one can easily move the weak ϵ-net points so that they still form weak ϵ-net and furthermore each of the weak ϵ-net points lies on the intersection of two lines, each of which is defined by two points in P. Similarly in $\mathbb{R}^d$, it is possible to move the points to the intersection of d hyperplanes, each of which is defined by d points of P. Therefore, any of the points is a *product* of d^2 points. This observation implies that if there is a small weak ϵ-net, it should be possible to construct it from a small number of the input points. However, all known constructions require $\Omega(\epsilon^{-d})$ points. We show that it is indeed possible to construct a weak ϵ-net from a random sample of size $O(\epsilon^{-1}\log\epsilon^{-1})$. Our algorithm first constructs a strong ϵ-net with of size $O(\epsilon^{-1}\log\epsilon^{-1})$ for a range space of a finite VC dimension and then takes certain *products* over it to produce the weak ϵ-net with respect to convex ranges. Apart from giving strength to the conjecture of Matousek and Wagner [37], the proof reveals an interesting connection between strong and weak ϵ-nets and shows that random sampling can be used to construct weak ϵ-nets. It also shows a connection between the Hadwiger-Debrunner (p, q) theorem and weak ϵ-nets.

Chapter 2

Strong ϵ-Nets

In Chapter 1, we mentioned the Strong ϵ-net theorem (Theorem 1) which states that any range space of a finite VC dimension d admits a strong ϵ-net of size $O(\frac{d}{\epsilon}\log\frac{d}{\epsilon})$. However, many range spaces, typically range spaces arising in geometry, admit strong ϵ-nets of size $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$. For, example, Pach and Woeginger [44] proved that halfspaces and translates of polytopes in $\mathbb{R}^2$ admit strong ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek et al. [36] proved that halfspaces in $\mathbb{R}^3$ and certain special families of pseudo-disks in $\mathbb{R}^2$ (they require that there is exactly one pseudo-disk through any three non-collinear points in the plane) admit strong ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek later found a shorter proof for the existence of $O(\frac{1}{\epsilon})$ size strong ϵ-nets for halfspaces in $\mathbb{R}^3$ via shallow cuttings [34].

In this chapter, we first give a new construction of strong ϵ-nets of size $O(\frac{1}{\epsilon})$ for halfspaces in $\mathbb{R}^2$ which leads to fast algorithm for computing ϵ-nets for halfspaces in $\mathbb{R}^2$. We then describe a general techniques for proving $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ upper bounds on the size of strong ϵ-nets admitted by various range spaces. We then show how to construct the strong ϵ-net efficiently. In the following, we will just write ”ϵ-net” for ”strong ϵ-net”.

2.1 Geometric Range Spaces

Recall that a range space $\mathcal{R}$ is a pair (X, R) where X is a *ground set* (possibly infinite) and R is a set of subsets of X. The elements of R are called *ranges*. For any $Y \subseteq X$ we call $R|_Y = \{r \cap Y : r \in S\}$ the *projection* of R on Y. The projection of $\mathcal{R}$ on Y is the range space $\mathcal{R}|_Y = (Y, R|_Y)$. The range spaces that we consider in this chapter are induced by finite sets of points and geometric objects. Let P be a finite set of points and S a finite set of geometric objects. For an object $s \in S$, let $P(s)$ be the set of points contained in s, i.e., $P(s) = \{p \in P : p \in s\}$. Similarly for a point p, let $S(p)$ be the set of objects containing p, i.e., $S(p) = \{s \in S : p \in s\}$. The sets P and S induce two natural kinds of range spaces. If we treat the set of points P as the ground set and the let the objects in S define the ranges, we get the range space $(P, \{P(s) : s \in S\})$ which we call the *primal range space induced by P and S* and denote it by $\mathcal{R}(P, S)$. On the other hand, if we think of the set of objects S as the ground set and let the points in P define the ranges, we get the *dual range space induced by P and S*, denoted by $\mathcal{R}^*(P, S) = (S, \{S(p) : p \in P\})$.

2.2 ϵ-Nets for halfspaces in $\mathbb{R}^2$

Let P be a finite set of n points in $\mathbb{R}^2$ and S be the set of all halfspaces in the $\mathbb{R}^2$. Pach and Woeginger [44] proved that the range space $\mathcal{R}(P, S)$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$. We give an alternate proof of this fact which allows us to compute an ϵ-net of the same size in $O(n \log \frac{1}{\epsilon})$ time.

For simplicity, let us assume that P is general position i.e. no three points of P are on the same line. Let p be a vertex of $\mathrm{CH}(P)$, the convex hull of P. Let $r_1, r_2, \cdots, r_k$ be rays

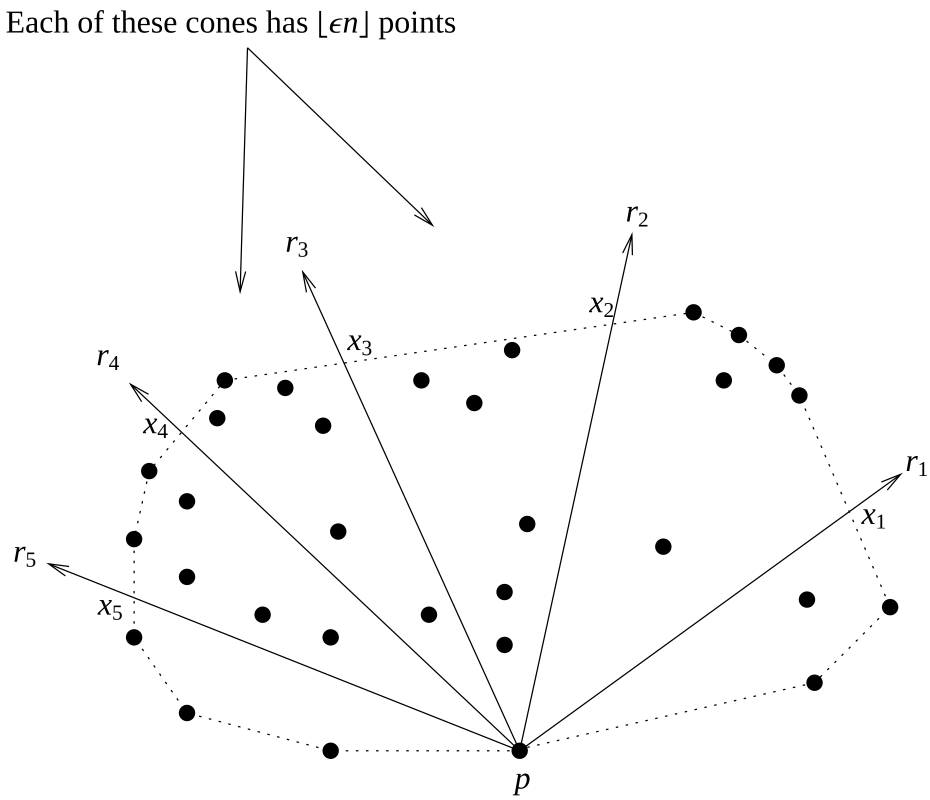

Figure 2.1: Construction of ϵ-net for halfspaces in $\mathbb{R}^2$.

emanating from p, where $k = \lceil \frac{1}{\epsilon} \rceil$, so that the cone defined by any two consecutive rays contains at most ϵn points of P (see Figure 2.1). For each ray r_i, let e_i be the edge of the convex hull of P that intersects r_i. Let u_i and v_i be the end points of e_i and x_i the intersection point of r_i and e_i. Let $Y = \bigcup_i \{u_i, v_i\}$.

Claim 1. *$E = Y \cup \{p\}$ is an ϵ-net for $\mathcal{R}(P, S)$. $|E| = O(\frac{1}{\epsilon})$.*

Proof. Consider a halfspace h that contains more than ϵn points from P. Assume that h contains points from a cone c defined by rays r_i and r_{i+1}. Then, h either contains one of the points x_i, x_{i+1}, p or it doesn't contain points from any other cone. In the latter case h contains at most ϵn points from P since c contains at most ϵn points from P. If h contains

x_i then it contains either u_i or v_i and hence is hit by E. Similarly, it is hit by E if it contains x_{i+1}. If h contains p, it is again hit by E since $p \in E$. By construction, $|E|$ is at most $1 + 2\lceil \frac{n}{\lfloor \epsilon n \rfloor} \rceil$. For large enough n this is arbitrarily close to $1 + 2\lceil \frac{1}{\epsilon} \rceil$. $\square$

We now show that such an ϵ-net can be constructed in $O(n \log \frac{1}{\epsilon})$ time.

Theorem 5. *Given a set P of n points in the plane, an ϵ-net of size $2\lceil \frac{1}{\epsilon} \rceil + 1$ with respect to halfspaces can be constructed in $O(n \log \frac{1}{\epsilon})$ time.*

Proof. Let p be a vertex of $\mathrm{CH}(P)$. Such a point is computed in $O(n)$ time by picking the lexicographically smallest point of P. Let $r_1, r_2, \cdots, r_k$ be rays emanating from p as before. In order to construct the ϵ-net described in Claim 1, we just need to compute these rays and the edge e_i of $\mathrm{CH}(P)$ intersecting each ray r_i. If $k = 1$, then we pick the ray r_1 to be any ray passing through p that intersects the interior of some edge of $\mathrm{CH}(P)$. Such a ray can be computed in $O(n)$ time. We can then compute the edge of $\mathrm{CH}(P)$ intersecting r_1 in $O(n)$ time by using the algorithm used in [30] for computing the bridge intersecting a given ray. Otherwise, if $k > 1$, we first recursively compute the odd numbered rays and the edges of $\mathrm{CH}(P)$ intersecting them. The even numbered ray r_i is chosen such that it roughly bisects the set of points P_i lying in the cone c_i defined by the odd numbered rays r_{i-1} and r_{i+1}. This is done in $O(n_i)$ time, where $n_i = |P_i|$, using a median computation algorithm. To compute e_i, we use the observation that e_i is either identical to one of the edges e_{i-1} or e_{i+1} or both its end points are in c_i. We compute the edge e'_i of $CH(P_i \cup \{p\})$ in $O(n_i)$ time using the bridge computation algorithm of [30]. The edge e_i is the edge among e_{i-1}, e_{i+1} and e'_i whose intersection with r_i is the furthest from P. Since $\sum_{\text{even i}} n_i = n$, the total time take to compute the rays r_i and the edges e_i for even indices i is $O(n)$. Hence, the overall running time $T(k,n)$ of the algorithm is given by the equations

$$T(k,n) = T(\lceil k/2 \rceil, n) + O(n)$$
$$T(1,n) = O(n)$$

Hence, $T(k,n) = O(n \log k) = O(n \log \frac{1}{\epsilon})$.

$\square$

We now describe a general technique to prove the existence of ϵ-nets of size $o(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$. We first give a simple proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspace ranges in three dimensions and then extend the technique to range spaces which satisfy certain simple conditions. We also show how the existence proofs can be converted into efficient algorithms for computing small ϵ-nets.

2.3 Halfspaces in $\mathbb{R}^3$

In this section we give a simple proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspaces in $\mathbb{R}^3$. This result was first proved in [36] and later a simpler proof appeared in [34]. For convenience, we consider the ϵ-net problem for the dual range space induced by finite sets of points and halfspaces in $\mathbb{R}^3$. The existence of ϵ-nets of size $O(\frac{1}{\epsilon})$ for such dual range spaces implies the same for the primal range spaces since the roles of points and halfspaces can be exchanged by using projective duality (see [18]).

The dual range space induced by a set of points and a set of halfspaces in $\mathbb{R}^3$ has the halfspaces as the ground set and each point defines a range which is the set of halfspaces containing that point. This range space clearly has a finite scaffold dimension and hence a finite VC dimension. Therefore, it follows from the ϵ-net theorem (Theorem 1) that such a range space admits an ϵ-net whose size depends only on the parameter ϵ. In other words, when ϵ is a constant, there is an ϵ-net of constant size. We use this fact and the following claim to prove that the dual range space induced by a given finite set P of points and a set H of halfspaces in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.

Claim 2. *Given any finite set Q of points in $\mathbb{R}^3$, there exists a graph $G_Q = (Q, E_Q)$ with at most $4|Q|$ edges such that for any halfspace h in $\mathbb{R}^3$, the subgraph of G_Q induced by the points of Q contained in h (i.e. $Q \cap h$) is connected.*

Proof. We construct G_Q as follows: Let $Q' \subseteq Q$ be the vertices of the convex hull $\mathrm{CH}(Q)$ of Q. We include the edges of the 1-skeleton of $\mathrm{CH}(Q)$ (i.e. the graph with the vertices of $\mathrm{CH}(Q)$ as the vertex set and the edges (1-faces) as the edge set) in G_Q. For each of the points $q \in Q \backslash Q'$, we pick a tetrahedron containing q whose vertices are in Q' (there is always such a tetrahedron by Carathéodory's theorem [35]) and put edges between q and each of the four corners of this tetrahedron. The construction of G_Q is complete. It contains at most $4|Q|$ edges since the 1-skeleton of $\mathrm{CH}(Q)$ is a planar graph and each point in $Q \backslash Q'$ has degree four. For any halfspace $h \in \mathbb{R}^3$, the subgraph of G_Q induced by the points in $Q' \cap h$ is obviously connected and each point in $(Q \backslash Q') \cap h$ is connected to at least one of the points of $Q' \cap h$. Therefore, the subgraph induced by the points in $Q \cap h$ is connected. □

For the range space $\mathcal{R}^*(P, H)$ defined by a set of points P and a set of H of n halfspaces in $\mathbb{R}^3$, call point $p \in P$ *heavy* if it is covered (contained) by more than ϵn halfspaces in H. Call a heavy point *moderately heavy* if the number of halfspaces covering it lies in the range $(\epsilon n, 2\epsilon n]$ and call it *very heavy* otherwise. We call a subset $Y \subseteq P$ a *moderate* ϵ-net for $\mathcal{R}^*(P, H)$ if for each moderately heavy point $p \in P$, a halfspace in Y covers p.

Claim 3. *The range space $\mathcal{R}^*(P, H)$ admits a moderate ϵ-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $M \subseteq P$ be the set of moderately heavy points in P. For each $p \in M$, let $H(p)$ denote the set of halfspaces in H which contain p. We say that two points $p, q \in M$ are *independent* if $|H(p) \cap H(q)| \leq \epsilon n/8$. Let $I \subseteq M$ be an inclusion-maximal set of pairwise independent points in M. The maximality of I implies that for any $p \in M$, there is a point $q \in I$, not necessarily different from p, such that $|H(p) \cap H(q)| > \epsilon n/8$. Since q is a moderately heavy point, $H(q) \leq 2\epsilon n$ and hence $|H(p) \cap H(q)| > |H(q)|/16$. This means that a $\frac{1}{16}$-net for $\mathcal{R}^*(P, H(q))$ hits p. In other words, $Z = \bigcup_{r \in I} Y_r$, where Y_r denotes a $\frac{1}{16}$-net for $\mathcal{R}^*(P, H(r))$, is a moderate ϵ-net for $\mathcal{R}^*(P, H)$. The size of such a net is $O(t)$, where $t = |I|$, since each of the $\frac{1}{16}$-nets is of constant size. We now show that $t = O(\frac{1}{\epsilon})$.

By Claim 2, there is a graph $G_I = (I, E_I)$ such that $|E_I| \leq 4t$ and for any $h \in H$, $|h \cap I|$ induces a connected subgraph of G_I. We say that a halfspace h contains an edge $e \in E_I$ if both the endpoints of e are contained in h. We denote the set of halfspaces containing an edge e by $H(e)$. For any halfspace h, let $n_h = |h \cap I|$ and let m_h be the number of edges contained in h. Since $h \cap I$ induces a connected subgraph of G_I, $n_h - m_h \leq 1$ for each h. Summing over the n halfspaces in H,

$$\sum_{h \in H} n_h - \sum_{h \in H} m_h \leq n. \tag{2.1}$$

Now, since each $p \in I$ is a heavy point,

$$\sum_{h \in H} n_h = \sum_{p \in I} |H(p)| > t\epsilon n. \tag{2.2}$$

Each edge in E_I is contained in at most $\epsilon n/8$ halfspaces since both its endpoints belong to I. Therefore,

$$\sum_{h \in H} m_h = \sum_{e \in E_I} |H(e)| \leq |E_I| \frac{\epsilon n}{8} \leq 4t \frac{\epsilon n}{8}. \tag{2.3}$$

It follows from (2.1),(2.2) and (2.3) that $t \leq 2/\epsilon$ and hence Z is a moderate ϵ-net of size $O(\frac{1}{\epsilon})$. □

Theorem 6. *The dual range space $\mathcal{R}^*(P,H)$ induced by a set of points P and a set of halfspaces H in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $M \subseteq H$ be a 2ϵ-net for $\mathcal{R}^*(P,H)$ and Z a moderate ϵ-net for $\mathcal{R}^*(P,H)$. Then $Z \cup M$ is an ϵ-net for $\mathcal{R}^*(P,H)$ since Z covers all of the moderately heavy points and M covers all the very heavy points. By Claim 3, there exists a moderate ϵ-net of size $O(1/\epsilon)$. If we denote the size of the smallest ϵ-net admitted by $\mathcal{R}^*(P,H)$ by $f(\epsilon)$, we have

$$f(x) = 0, \quad \forall\, x \geq 1,$$
$$f(\epsilon) \leq O(\frac{1}{\epsilon}) + f(2\epsilon).$$

It follows that $f(\epsilon) = O(\frac{1}{\epsilon})$. □

Using projective duality between points and halfspaces, we also obtain the next theorem.

Theorem 7. *The primal range space $\mathcal{R}(P,H)$ induced by a finite set of points P and a finite set of halfspaces H in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

2.4 Abstract Framework

The proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for range spaces induced by points and halfspaces in R^3 can be adapted to any range space $\mathcal{R} = (X, S)$ that has the properties that we have exploited in the proof. In the following, we denote the set of ranges containing a particular element $x \in X$ by $S(x)$.

Theorem 8. *Any range space $\mathcal{R} = (X, S)$ satisfying the following two conditions admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

1. *For any $0 < \epsilon < 1$ and for any $Y \subseteq X$, $\mathcal{R}|_Y$ admits an ϵ-net whose size depends only on ϵ.*

2. *There exist constants $\alpha > 0, \beta \geq 0$ and $\tau > 0$ s.t. for any $I \subseteq S$, there is a graph $G_I = (I, E_I)$ with $|E_I| \leq \beta|I|$ so that for any element $x \in X$ we have $m_x \geq \alpha n_x - \tau$, where $n_x = |I(x)|$ and m_x is the number of edges in E_I whose both endpoints (which are ranges) contain x.*

Proof. The proof is analogous to the proof for halfspaces in $\mathbb{R}^3$. As before, we call a range $s \in S$ *heavy* if $|s| > \epsilon n$, where $n = |X|$. We say that a heavy range s is *very heavy* if $|s| > 2\epsilon n$ and *moderately heavy* otherwise. An ϵ-net for $\mathcal{R}$ is a subset $Y \subseteq X$ which *hits* all heavy ranges i.e. each heavy range contains at least one element of Y. A moderate ϵ-net is a subset $Z \subseteq X$ which hits all moderately heavy ranges.

We show that $\mathcal{R}$ admits a moderate ϵ-net of size $O(\frac{1}{\epsilon})$. From this we can conclude the existence of an $O(\frac{1}{\epsilon})$ size ϵ-net for $\mathcal{R}$ by an argument analogous to the proof of Theorem 6. We say that two ranges s and s' are *independent* if $|s \cap s'| \leq \frac{\alpha}{2\beta}\epsilon n$. Let $I \subseteq S$ be an inclusion-maximal set of pairwise independent moderately heavy ranges. Then for each $s' \in S$, there is a set $s \in I$ such that $|s \cap s'| > \frac{\alpha}{2\beta}\epsilon n$ which implies that $|s \cap s'| > \frac{\alpha}{4\beta}|s|$ since $|s| \leq 2\epsilon n$. Therefore, an $\frac{\alpha}{4\beta}$-net for $\mathcal{R}_s$ hits s'. This means that $Z = \bigcup_{r \in I} Y_r$, where Y_r denotes a $\frac{\alpha}{4\beta}$-net for $\mathcal{R}|_r$, is a moderate ϵ-net for $\mathcal{R}$. Moreover, $|Z| = O(t)$ where $t = |I|$, since each of the $\frac{\alpha}{4\beta}$-nets are of constant size due to the first condition in the statement of the theorem. Now we show that $t = O(\frac{1}{\epsilon})$.

Let $G_I = (I, E_I)$ be the graph ensured by the second condition in the statement of the theorem. For an edge $e = (s, s') \in E_I$, let $X(e) = |s \cap s'|$. Since the ranges in I are pairwise independent, $X(e) \leq \frac{\alpha}{2\beta}\epsilon n$. For each $x \in X$, we have $\alpha n_x - m_x \leq \tau$, where $n_x = |I(x)|$ and m_x is the number of edges in E_I whose both endpoints contain x. Summing over all $x \in X$, we have:

$$\sum_{x\in X} \alpha n_x - \sum_{x \in X} m_x \leq \tau n. \tag{2.4}$$

Now,

$$\sum_{x\in X} n_x = \sum_{s\in I} |s| \geq t\epsilon n \tag{2.5}$$

since each $s \in I$ is heavy. Also, since $X(e) \leq \frac{\alpha}{2\beta}\epsilon n$ for each $e \in E_I$ and $|E_I| \leq \beta t$,

$$\sum_{x\in X} m_x = \sum_{e\in E_I} X(e) \leq |E_I|\frac{\alpha}{2\beta}\epsilon n \leq \beta t \frac{\alpha}{2\beta}\epsilon n = \frac{\alpha}{2} t\epsilon n. \tag{2.6}$$

From (2.4), (2.5) and (2.6) we get:

$$\alpha t\epsilon n - \frac{\alpha}{2} t\epsilon n \leq \tau n \implies t \leq \frac{2\tau}{\alpha} \cdot \frac{1}{\epsilon}. \tag{2.7}$$

Since α and τ are constants, $t = O(\frac{1}{\epsilon})$. Hence, Z is a moderate ϵ-net of size $O(\frac{1}{\epsilon})$ for $\mathcal{R}$ and we conclude from a calculation similar to the one in the proof of Theorem 6 that $\mathcal{R}$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$. $\square$

The second condition of Theorem 8 requires $|E_I|$ to be $O(|I|)$. It is natural to expect that if instead we had $|E_I| \leq |I|b(|I|)$ for some *small* function $b(\cdot)$ then we should still be able to prove the existence of a correspondingly *small* ϵ-net. Indeed, this is true and the following theorem can be proved along the lines of the proof of Theorem 8.

Theorem 9. *Let $\mathcal{R} = (X, S)$ be a range space satisfying the following two conditions:*

1. *For any ϵ and for any $Y \subseteq X$, $\mathcal{R}|_Y$ admits an ϵ-net whose size depends only on ϵ.*

2. *There exist constants $\alpha > 0, \tau > 0$ and a positive non-decreasing sublinear function $b(\cdot)$ s.t. for any $I \subseteq S$, there is a graph $G_I = (I, E_I)$ with $|E_I| \leq |I|\, b(|I|)$ so that for any element $x \in X$ we have $m_x \geq \alpha n_x - \tau$, where $n_x = |I(x)|$ and m_x is the number of edges in E_I whose both endpoints (which are ranges) contain x.*

Then, $\mathcal{R}$ admits an ϵ-net of size $O\left(\frac{1}{\epsilon} \cdot (4\tau)^{f^*\left(\frac{2}{\epsilon}\right)} \tilde{f}\left(\frac{2}{\epsilon}\right)\right)$ where $f^*(\cdot)$ and $\tilde{f}(\cdot)$ are defined as:

$$f^*(k) = \begin{cases} 0, & \text{if } f(k) \geq k \\ 1 + f^*(f(k)), & \text{otherwise} \end{cases}$$

$$\tilde{f}(k) = \begin{cases} 1, & \text{if } f(k) \geq k \\ f(k) \cdot \tilde{f}(f(k)), & \text{otherwise} \end{cases}$$

The proof Theorem 8 can be easily adapted for the case in which the vertices have positive weights (instead of all vertices having weight 1). The same holds for Theorem 9.

2.4.1 Algorithmic Issues

The proof of Theorem 8 suggests the following simple algorithm for computing an ϵ-net for a range space $\mathcal{R} = (X, S)$ satisfying the conditions of the theorem: Start with an empty set as the ϵ-net and look at the ranges in S one by one. Let s be the range currently being considered. If s is already hit by the ϵ-net we have built so far then we ignore it. If s is large ($|s| > \epsilon n$) and is not already hit we compute an $\frac{\alpha}{4\beta}$-net for $\mathcal{R}|_s$ and add it to the current ϵ-net. The pseudocode is shown in Algorithm 1.

Algorithm: Compute ϵ-net

Input: A range space $\mathcal{R} = (X, S)$

Output: An ϵ-net N for $\mathcal{R}$

$N = \emptyset$ `// The` ϵ`-net is initially empty`

forall $(s \in S)$ **do**

    if $|s| > \epsilon n$ *and* $s \cap N = \emptyset$ **then**

        Pick an $\frac{\alpha}{4\beta}$-net M_s for $\mathcal{R}|_s$

        Set $N := N \cup M_s$

    end

end

return N

Algorithm 1: Compute ϵ-net

The set N constructed in Algorithm 1 is an ϵ-net by construction. We still need to argue that it has a small size. Consider the subset of ranges $S' \subseteq S$ whose sizes are in the interval $(\epsilon n, 2\epsilon n]$ and for which Line 5 is executed. The ranges in S' form an independent set i.e. for any two ranges $s, s' \in S'$, $|s \cap s'| \leq \frac{\alpha}{2\beta}\epsilon n$ and hence by the argument used in the proof of Theorem 8, $|S'| \leq \frac{2\tau}{\alpha}\frac{1}{\epsilon}$. Similarly, the number of ranges whose sizes are in the interval $(2^k\epsilon n, 2^{k+1}\epsilon n]$ and for which Line 5 is executed is at most $\frac{2\tau}{\alpha}\frac{1}{2^k\epsilon}$ for any integer $k \geq 0$. Hence, the total number of ranges for which Line 5 is executed is $O(\frac{1}{\epsilon})$. Since the size of each $\frac{\alpha}{4\beta}$-net computed in Line 5 is a constant, the size of the ϵ-net computed is $O(\frac{1}{\epsilon})$.

The implementation of Line 5 depends on the range space under consideration. Assuming that it takes constant time to execute Line 5 and to check whether an element of the ground set belongs to a given range, the overall running time of the algorithm is $O(mn)$. The size of each range is computed by checking for each element of the ground set whether it belongs to the range. Checking whether a certain range is hit by the current ϵ-net is again done by checking whether any of the elements of the current ϵ-net is contained in the range. This takes $O(\frac{1}{\epsilon})$ time. We can assume that $\epsilon > \frac{1}{n}$ since otherwise we can pick the whole ground set as the ϵ-net. Therefore the total time taken is $O(mn)$.

If $\mathcal{R}$ has a finite VC dimension d, then we can compute the $\frac{\alpha}{4\beta}$-net required in Line 5 of Algorithm 1 by random sampling. By the ϵ-net theorem (Theorem 1), a random sample of size $O(\frac{d}{\delta}\log\frac{d}{\delta})$, where $\delta = \frac{\alpha}{4\beta}$, is an $\frac{\alpha}{4\beta}$-net with high probability. Suppose that the probability is more than $\frac{1}{2}$. Then, in expectation in at least half of the cases the random sample is an $\frac{\alpha}{4\beta}$-net. If we let S' be the subset of ranges $S' \subseteq S$ whose sizes are in the interval $(\epsilon n, 2\epsilon n]$ and for which Line 5 is executed successfully (i.e. we get a correct $\frac{\alpha}{4\beta}$-net), then as before, it can be argued that $|S'| \leq \frac{2\tau}{\alpha}\frac{1}{\epsilon}$. Therefore, the total number of ranges whose sizes are in the interval $(\epsilon n, 2\epsilon n]$ and for which Line 5 is executed (either successfully or unsuccessfully) is at most $2\frac{2\tau}{\alpha}\frac{1}{\epsilon}$ in expectation. As before, we can argue that the total number of times Line 5 is executed is $O(\frac{1}{\epsilon})$. Hence the size of the ϵ-net computed is $O(\frac{1}{\epsilon})$ and the running time of the algorithm is $O(mn)$.

The running time $O(mn)$ can often be prohibitive since m can be quite large. In such cases, a standard technique is to first construct $\frac{\epsilon}{2}$-approximation A for $\mathcal{R}$ and then use a subspace oracle to enumerate the hyperedges of $\mathcal{R}|_A$. An $\frac{\epsilon}{2}$-net for $\mathcal{R}|_A$ is an ϵ-net for $\mathcal{R}$. If the scaffold dimension of $\mathcal{R}$ is d, then a subspace oracle enumerates the distinct ranges of $\mathcal{R}|_A$ in $O(|A|^{d+1})$ time. The $\frac{\epsilon}{2}$-approximation A of size $O(\epsilon^{-2}\log\epsilon^{-1})$ can be computed deterministically in $O(n\epsilon^{-2d}\log^d\epsilon^{-1})$ time [39]. The distinct ranges in $\mathcal{R}|_A$ can then be enumerated in $O((\epsilon^{-2}\log\epsilon^{-1})^{d+1}) = O(n\epsilon^{-2d}\log^d\epsilon^{-1})$ time.[1] Now since we have only $O((\epsilon^{-2}\log\epsilon^{-1})^d)$ ranges to deal with, we can use the previous $O(mn)$ time algorithm and the overall running time of our algorithm remains $O(n\epsilon^{-2d}\log^d\epsilon^{-1})$.

2.5 Geometric Applications

In the following, we present several applications of Theorem 8. The geometric range spaces that we consider here have finite VC dimension and hence automatically satisfy the first condition of Theorem 8 (due to Theorem 1). Hence, we only prove that they satisfy the second condition and conclude the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for them. Most of the results presented here have been proved before using various techniques. Apart from Theorems 14 and 16, which were not previously known, the remaining results also follow from the framework of Clarkson and Varadarajan [16]. We include them here in order to demonstrate that they follow from our framework too. Also, in some cases, our technique

[1] again assuming that ϵ is a constant

leads to a simpler proof. Theorems 14 and 16 show a way to overcome the limitations of the technique used in [16]. Most of the definitions given in this section are more thoroughly explained in [2].

2.5.1 Translates of Orthants in $\mathbb{R}^3$

We will show that the dual range space induced by a finite set of points and translates of an orthant in $\mathbb{R}^3$ (also called an octant) admits an ϵ-net of size $O(\frac{1}{\epsilon})$. As mentioned earlier, this result also follows from the framework of Clarkson and Varadarajan [16].

Let P be a finite set of points in $\mathbb{R}^3$ and let O be the set of all translates of some orthant in $\mathbb{R}^3$. We will also denote a point $p \in \mathbb{R}^3$ as (x_p, y_p, z_p), where x_p, y_p and z_p are the x, y and z coordinates of p. For $p, q \in \mathbb{R}^3$, we will write $p \geq q$ iff $x_p \geq x_q, y_p \geq y_q$ and $z_p \geq z_q$. We define the notation $p \leq q$ in a similar manner. W.l.o.g. we assume that the orthants in O are axis-parallel, and every $T \in O$ is of the form $\{(x, y, z) : x \geq x_T, y \geq y_T, z \geq z_T\}$, for some $(x_T, y_T, z_T) \in \mathbb{R}^3$, which we call the *corner* of T. For any $T \in O$ and any $Q \subseteq P$, let $Q(T) = Q \cap T$. Note that if an orthant $T \in O$ contains some $p \in Q$, then it also contains any point $q \in Q$ such that $q \geq p$. For any two points $p, q \in Q$, we define the *minimal common orthant of p and q, $T_{p,q} \in O$,* as the minimal (w.r.t. inclusion) orthant that contains both p and q. The corner of $T_{p,q}$ is the point $(\min\{x_p, x_q\}, \min\{y_p, y_q\}, \min\{z_p, z_q\})$.

Lemma 1. *For any $Q \subseteq P$, there is a graph $G_Q = (Q, E_Q)$, such that $|E_Q| \leq 3|Q|$ and for any $T \in O$, there are at least $\frac{1}{2}|Q(T)| - 1$ edges among the points in $Q(T)$.*

Proof. For a point $p \in Q$ we define the *x-neighbor* of p as $N_x(p) = \arg\max_q\{x_q : y_q \geq y_p, z_q \geq z_p, q \in Q \setminus \{p\}\}$. The y- and z-neighbors $N_y(p)$ and $N_z(p)$ are defined analogously. Note that it is not necessary that all three of $N_x(p), N_y(p), N_z(p)$ exist for all

$p \in Q$. For every $p \in Q$, we add to E_Q the edges $(p, N_x(p)), (p, N_y(p))$ and $(p, N_z(p))$, whenever $N_x(p), N_y(p)$ and $N_z(p)$ exist, respectively. The construction of G_Q is now complete. Clearly, $|E_Q| \leq 3|Q|$, since every point in Q accounts for at most 3 edges.

Consider any orthant $T \in O$. We claim that there is at most one point $p \in Q(T)$ that does not share an edge with another point in $Q(T)$. This will immediately imply that the number of edges whose both endpoints are contained in Q is at least $\frac{1}{2}|Q(T)| - 1$. Assume, for contradiction, that there are $p, q \in Q(T)$, $p \neq q$, neither of which shares an edge with another point in $Q(T)$. Since $p, q \in T$, their minimal common orthant $T_{p,q}$ is also contained in T. W.l.o.g., assume that the corner of $T_{p,q}$ is the point $o_{p,q} = (x_p, y_p, \min\{z_p, z_q\})$, i.e. at least two coordinates of the corner (namely the x and y coordinates) are equal to the corresponding coordinates of p (any other case can be treated similarly). Consider the set of points $Z_p = \{q' \in Q \setminus \{p\} : x_{q'} \geq x_p, y_{q'} \geq y_p\}$. The z-neighbor of p is given by $N_z(p) = \arg\max_{q'}\{z_{q'} : q' \in Z_p\}$. Since $q \in Z_p(T)$, it must be that $z_q \leq z_{N_z(p)}$, implying that $N_z(p) \geq o_{p,q}$, i.e. $N_z(p) \in T$. But then the edge $(p, N_z(p)) \in E_Q$ contradicting our assumption that p does not share an edge with another point in $Q(T)$. Therefore, there is at most one point in $Q(T)$ which does not share an edge with another point in $Q(T)$, thus proving the claim. □

Lemma 1 and Theorem 8 imply the following theorem:

Theorem 10. *The dual range space defined by a finite set of points and a finite set of translates of an orthant in* $\mathbb{R}^3$ *admits an* ϵ*-net of size* $O(\frac{1}{\epsilon})$.

The above theorem also implies the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for the primal range space. To see this, note that we can substitute every orthant T by its corner o_T, and every point $p = (x_p, y_p, z_p)$ by an orthant of the form $\{(x, y, z) : x \leq x_p, y \leq y_p, z \leq z_p\}$. This

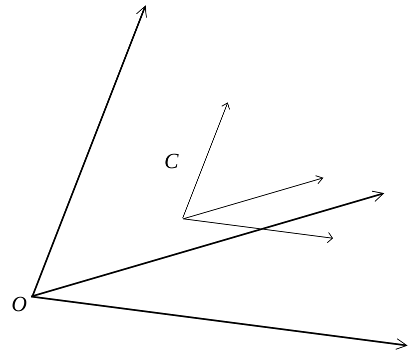

Figure 2.2: Cone C behaves like an orthant in the oblique coordinate system with origin O.

preserves the incidences between orthants and points. Therefore, we have:

Theorem 11. *The primal range space defined by a finite set of points and a finite set of translates of an orthants in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Translates of Polytopes in $\mathbb{R}^3$: We proved that the range spaces induced by a set of n points and translates of an orthant in $\mathbb{R}^3$ admit a small ϵ-net. It is possible to prove the same result for the range space induced by translates of a polytope in $\mathbb{R}^3$ with a finite number of vertices. We give a brief sketch here. Suppose that instead of translates of an orthant, we had translates of a cone with a triangular cross section i.e., a cone which is the intersection of three halfspaces passing through a point. This situation is not very different since cones with a triangular cross-section behave like orthants in a suitable oblique coordinate system (see Fig. 2.2). In other words, it is easy to apply an affine transform to make the translates of the cone look like orthants while preserving the incidences between the points and the translates of the cone. Therefore, the range space induced by a set of points and translates of such *triangular* cones in $\mathbb{R}^3$ also admits an ϵ-net of size $O(1/\epsilon)$. From this, it is easy to show that the range space induced by translates of a tetrahedron Δ admits an ϵ-net of size $O(1/\epsilon)$. We only need to consider a fine enough (oblique) gridding of $\mathbb{R}^3$ so that every

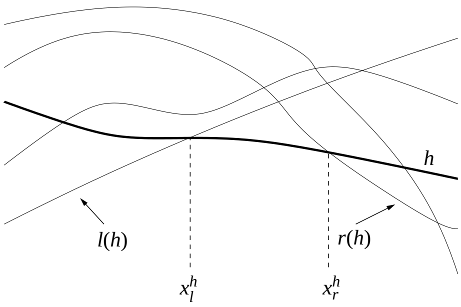

Figure 2.3: $l(h)$ and $r(h)$ for the pseudo-halfplane h in a family of pseudo-halfplanes.

translate of Δ has at most one corner in any cell and at the same time it does not intersect more than a constant number t of cells. Then, inside any cell, we can treat the translates of Δ intersecting it as cones. There are four different types of cones corresponding to the four corners of Δ. If the cell under consideration has n_i points in it, then we construct an ϵ'-net, where $\epsilon' = \epsilon \frac{n}{tn_i}$ of size $O(n_i/\epsilon n)$ with respect to translates of each of the four different kinds of cones. It is not hard to argue that the union of these ϵ'-nets for the cells gives an ϵ-net for the range space induced by translates of Δ and total size is still $O(1/\epsilon)$ since $\sum_{cells} n_i = n$. Since a convex polytope in $\mathbb{R}^3$ with k vertices can be triangulated with $O(k)$ tetrahedra, it follows that translates of a convex polytope in $\mathbb{R}^3$ with k vertices admit an ϵ-net of size $O(\frac{k^2}{\epsilon})$. We just need to construct an ϵ/k-net with respect to the range spaces induced by the translates of each of the $O(k)$ tetrahedra used in the triangulation of the given polytope.

2.5.2 Pseudo-Halfplanes in $\mathbb{R}^2$

A family of (x-monotone) *pseudo-lines* in the plane is a set of graphs of continuous univariate functions, that intersect in at most one point and cross at that point. A family of *pseudo-halfplanes* is a set of closed sets in the plane whose boundaries form a family of pseudo-lines. For convenience, we will just write *halfplanes* for pseudo-halfplanes. For any halfplane h, we will denote the function tracing its boundary by f_h. With a slight abuse of notation, we will also refer to the boundary of h by f_h. A family of *upper* halfplanes is a set of halfplanes each of which is bounded from below, i.e. for each halfplane h in the set, $h = \{(x, y) \in \mathbb{R}^2 : y \geq f_h(x)\}$. Similarly, a family of *lower* halfplanes is a set of halfplanes each of which is bounded from above.

Let $\mathscr{H}$ be a finite family of upper halfplanes and $X \subseteq \mathscr{H}$. For any $p \in \mathbb{R}^2$, we will denote the set of halfplanes in X that contain p by $X(p)$, i.e., $X(p) = \{h \in X : p \in h\}$.

Lemma 2. *For any $X \subseteq \mathscr{H}$, there is a graph $G_X = (X, E_X)$, such that $|E_X| \leq 2|X|$ and for any $p \in \mathbb{R}^2$, the halfplanes in $X(p)$ induce a connected subgraph (which therefore contains at least $|X(p)| - 1$ edges).*

Proof. For halfplanes h and h', we say that h lies *below* h' at x, if $f_{h'}(x) < f_h(x)$. Note that if a point $p = (x, y)$ is contained in some $h \in X$, then p is also contained in every h' lying below h at x. Therefore, in order to construct G_X in such a way that the set of halfplanes containing p induce a connected subgraph, it suffices to ensure that for all $x \in \mathbb{R}$, each halfplane $h \in X$ shares an edge with some halfplane $h' \in X$ below it at x (if one such h' exists). For simplicity, we assume that the boundaries f_h and $f_{h'}$, of any pair of halfplanes $h, h' \in X$, cross exactly once. (If there are h, h' such that the boundaries f_h and $f_{h'}$ never cross, then for all $x \in \mathbb{R}$, one of them, say h, lies above the other. Therefore, we can put an edge between h and h' and ignore pseudo-halfplane h.) For a halfplane $h \in X$, let $l(h)$ be the halfplane in X which lies below h for the maximal interval (w.r.t. inclusion) of the form $(-\infty, x)$. Similarly let $r(h)$ be the halfplane of X which lies below h for the maximal interval of the form $(x, +\infty)$. More formally,

$$\begin{aligned} l(h) &= \arg\max_{h' \in X} \{\, x_l : \quad f_{h'}(x_l) = f_h(x_l) \text{ and } \forall x < x_l,\ f_{h'}(x) < f_h(x) \,\}, \\ r(h) &= \arg\min_{h' \in X} \{\, x_r : \quad f_{h'}(x_r) = f_h(x_r) \text{ and } \forall x > x_r,\ f_{h'}(x) < f_h(x) \}. \end{aligned}$$

Let x_h^l and x_h^r be the x-coordinates of the intersection of f_h with $f_{l(h)}$ and $f_{r(h)}$, respectively.[2] For an example, see Fig. 2.3. It is easy to see that connecting every $h \in X$ with $l(h)$ and $r(h)$ gives the required graph. For all $x \in (-\infty, x_l^h)$, h is connected to the halfplane $l(h)$ lying below it and for all $x \in (x_r^h, +\infty)$, it is connected to the halfplane $r(h)$ lying below it. For any $x \in [x_h^l, x_h^r]$, there is no halfplane lying below h. □

It follows from Theorem 8 and Lemma 2 that the primal range space defined by a family of upper (or lower) pseudo-halfplanes and a set of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.

Theorem 12. *The primal range space induced by a finite family of pseudo-halfplanes $\mathscr{H}$ and a finite set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $\mathscr{H}_l$ and $\mathscr{H}_u$ be the sets of lower and upper halfplanes (respectively) in $\mathscr{H}$. We construct separate ϵ-nets N_l and N_u for $\mathcal{H}(P, \mathscr{H}_l)$ and $\mathcal{H}(P, \mathscr{H}_u)$. Then, $N_l \cup N_u$ gives an ϵ-net of size $O(\frac{1}{\epsilon})$ for $\mathcal{H}(P, \mathscr{H})$. □

Using the duality between a family of pseudo-lines and a set of points in the plane, as defined in [3], we can exchange the roles of points and upper halfplanes in Lemma 2 and prove the following:

[2] If $x_h^l > x_h^r$, we can completely ignore h for this stage, since the edges added to it in the first stage suffice. If x_h^l or x_h^r does not exist, we can consider it equal to $-\infty$ or $+\infty$, respectively.

Theorem 13. *The dual range space induced by a finite family of pseudo-halfplanes $\mathscr{H}$ and a finite set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

2.5.3 Pseudo-Parabolic Halfplanes in $\mathbb{R}^2$

A family of *pseudo-parabolas* is a set of graphs of continuous univariate functions every two of which intersect (and cross) in at most two points. (We assume that every tangency is equivalent to two intersections.) A family of *pseudo-parabolic halfplanes* (parabolic halfplanes for short) is a set of closed regions in the plane whose boundaries form a family of pseudo-parabolas. For a parabolic halfplane b, let f_b denote the function that defines the corresponding pseudo-parabola. We define *upper* and *lower* parabolic halfplanes just as we did for pseudo-halfplanes in Section 2.5.2.

Lemma 3. *Let $\mathscr{B}$ be a family of upper pseudo-parabolic halfplanes and P a set of points in $\mathbb{R}^2$. Then, for any $B \subseteq \mathscr{B}$ there is a graph $G_B = (B, E_B)$, such that $|E_B| \leq 6|B|$ edges, and for any $p \in P$ the parabolic halfplanes containing p induce a connected subgraph.*

Proof. We will assume that any $b_1, b_2 \in B$ cross exactly twice. It is not too hard to prove that for any $b_1, b_2 \in B$ that intersect only once, additional crossings can be created at the right of the rightmost point in P, without changing the incidences between the parabolic halfplanes and the points in P.[3] For each $b \in B$, we define $l(b), r(b), x_l^b, x_r^b$ just as in the proof of Lemma 2. Note that the definitions imply that x_l^b is the first point of intersection between b and $l(b)$, while x_r^b is the second point of intersection between $r(b)$ and b. We construct the edge set E_B in two stages. In the first stage, we connect every $b \in B$ with

[3] Again, if $b, b' \in B$ intersect tangentially or never intersect, then one of them, let it be b, lies completely above the other. Therefore, we can connect b to b' with an edge and ignore b from then on.

$l(b)$ and $r(b)$. This was we add at most $2|B|$ edges and every $b \in B$ is connected to another parabolic halfplane that lies below it, for all $x \in (-\infty, x_l^b) \cup (x_r^b, +\infty)$. In the second stage, for every $b \in B$, we restrict our attention to the interval $I_b = [x_b^l, x_b^r]$. Let σ_b denote the drawing of f_b restricted to I_b. We claim that the x-monotone curves σ_b (for all $b \in B$) form a set Γ of curves whose interiors do not cross. Assuming the contrary, say that the interiors of σ_{b_1} and σ_{b_2} intersect at a point (x_0, y_0). If this is the first intersection between the curves f_{b_1} and f_{b_2}, and for $x < x_0$, $f_{b_1}(x) < f_{b_2}(x)$, then $x_l^{b_2} \geq x_0$, contradicting the assumption that the interiors of σ_{b_1} and σ_{b_2} intersect at (x_0, y_0). Similarly, (x_0, y_0) cannot be the second point of intersection between the two curves either.

For each $b \in B$ and $x \in I_b$, we will ensure that b is connected to some $b' \in B$ such that $x \in I_{b'}$ and $\sigma_{b'}$ lies below σ_b at x (if such a b' exists). To achieve this, we consider a trapezoidal decomposition of the segments σ_b (as in [18], Chapter 6, noting that having curved segments instead of line segments doesn't cause any problems). The decomposition consists of at most $3|X| + 1$ non-overlapping trapezoids, whose upper and lower sides are parts of curves in Γ and the left and right sides are vertical line segments. For each trapezoid, we put an edge between the parabolic halfplanes corresponding to the upper and lower sides. The construction of G_B is now complete.

To see that for any $p \in \mathbb{R}^2$ the subgraph of G_B induced by the parabolic halfplanes in $B(p)$ is connected, observe first that the lower envelope of the arrangement of parabolic halfplanes in B is identical to the lower envelope of the segments in Γ. Now, for each $b \in B$ and $x \in \mathbb{R}$: If $x \in (-\infty, x_l^b) \cup (x_r^b, +\infty)$, then b is connected to some parabolic halfplane lying below it at x (if one such exists) due to the edges added in the first stage. Otherwise $x \in I_b$, and the only case in which b is not connected to any parabolic halfplane lying below it is when σ_b is the lowest segment of Γ at x. However, in that case, b is also the lowest parabolic halfplane of B at x. Therefore for every $x \in \mathbb{R}$, each parabolic halfplane $b \in B$ shares an edge with some $b' \in B$, such that $f_{b'}(x) < f_b(x)$, unless b is the lowest parabolic halfplane at x, implying thus that the subgraph is connected. Since the total number of edges added in the two stages is at most $5|B| + 1 \leq 6|B|$, the Lemma follows. □

The next theorem follows from Lemma 3 and Theorem 8.

Theorem 14. *The primal range space induced by a finite family of pseudo-parabolic halfplanes $\mathscr{B}$ and a set of points P in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Unfortunately, there is no duality (similar to the one used for pseudo-lines) known for a set of pseudo-parabolas and a set of points in $\mathbb{R}^2$. Therefore, a different technique is required in order to prove the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for the dual range space defined by a family of pseudo-parabolic halfplanes and a set of points in the plane. We will, in fact, derive it from a more general result proved in the next section (see Remark 1 at the end of the chapter).

2.5.4 k-admissible Regions in $\mathbb{R}^2$

Consider now a family of regions in $\mathbb{R}^2$, each of which is bounded by a closed Jordan curve. We will call it a family of *k-admissible* regions (for k even), if for any two s_1, s_2 of the regions, the Jordan curves bounding them cross [4] in $l \leq k$ points, (for some even l), and both $s_1 \setminus s_2$ and $s_2 \setminus s_1$ are connected regions. A family of 2-admissible regions is also called a family of *pseudo-disks*.

Let S be a family of k-admissible regions, and P a *finite* set of points in $\mathbb{R}^2$. For any $Q \subseteq P$, we will show that there is a plane multigraph G (a crossing-free drawing of planar

[4] two Jordan curves cross when at a certain point a curve passes from one side of another curve to the other

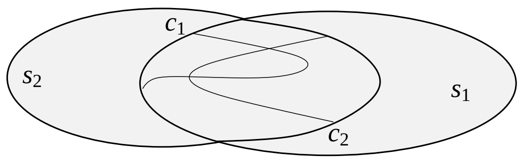

Figure 2.4: A chord of s_1 lies inside s_2 and vice-versa.

graph which may contain multiple edges between two vertices) with vertex set Q, such that the subgraph of G induced by the set of points contained in any of the regions $s \in S$, is connected. The graph will be given as the union of *connecting graphs* for each $s \in S$: For a region $s \in S$, we call a plane connected graph $G_s = (Q(s), E_s)$, where $Q(s) = Q \cap s$, an *s-connecting graph* if the drawing of the edges in E_s is strictly contained in s. Moreover, we say that a set of edges E *properly connects* a region $s \in S$, if there is a subset $E' \subseteq E$, such that the graph $G' = (Q(s), E')$ is an s-connecting graph.

In the following, whenever we refer to an edge, we will also refer to its drawing. We say that an edge *e pierces* a region s, if $s \setminus e$ has at least two connected components and not all the points of $Q \cap s$ lie in the same component. A *chord* of a region s is a Jordan arc with the endpoints lying on the boundary ∂s of s and the interior lying in the strict interior of s. If c is a chord of s, then $s \setminus c$ consists of exactly two connected regions.

Lemma 4. *Let $s_1, s_2 \in S$. Let c_1 be a chord of s_1 that lies in the interior of s_2 and c_2 be a chord of s_2 that lies in the interior of s_1. Then c_1 and c_2 cross at an even number of points.*

Proof. Fig. 2.4 shows a simple example which gives the intuition behind the lemma. We now prove it formally. Since c_1 is a chord of s_1, it splits s_1 into two parts A and B. The chord c_1 lies in the interior of s_2 and $s_1 \setminus s_2$ has at most one connected component. Therefore, exactly one of A, B is contained in the interior of s_2. (If both A, B lie in s_2 then s_1 does not contain any point of ∂s_2, and hence it cannot contain a chord of s_2 either.) Assume that

B is not contained in s_2. Since c_2 is a chord of s_2, the endpoints of c_2 are on ∂s_2. Also, A is contained in the strict interior of s_2 and therefore does not contain any points of ∂s_2. Hence, both endpoints of c_2 must lie in B, i.e. they are both on the same side of c_1 inside s_1. This immediately implies that c_1 and c_2 may only cross an even number of times. $\square$

The following lemma indicates how to construct the aforementioned plane multigraph.

Lemma 5. *For any $S' \subseteq S$ and any given set of pairwise non-crossing* "compulsory" *edges E_c, such that no edge $e \in E_c$ pierces any of the regions in S', there is a plane multigraph graph $G = (Q, \mathbb{E} \cup D)$, such that $E_c \cup D$ properly connects every $s \in S'$.*

Proof. Let $S' = \{s_1, s_2, \cdots, s_d\}$. We will use induction on the cardinality d of S'.

For $d = 1$, let I_1 be the set of points in the plane contained in the interior of s_1 and which do not lie in the interior of any of the edges in E_c. Since no edge in E_c pierces s_1, all points in $Q(s)$ belong to a connected component I'_1 of I_1. Therefore, there is a plane multigraph $G_1 = (Q, E_c \cup D_1)$, such that the edges in D_1 are strictly contained in I'_1 (and therefore in the strict interior of s_1) and $E_c \cup D_1$ properly connects s_1.

Assume now that for $d = l \geq 1$ and any compulsory set of pairwise non-crossing edges E_c which do not pierce any of $s \in S'$ there is a plane multigraph $G_l = (Q, E_c \cup E_l)$, with $E_c \cup E_l$ properly connecting $s_1, s_2, .., s_l$. For $d = l + 1$, let E'_l be the subset of edges in E_l that do not pierce s_{l+1}. Any edge $e \in E_l \setminus E'_l$ is split by ∂s_{l+1} into a set of segments. The segments that are contained in s and are not chords of s_{l+1} will be called *obstacles* (see for example Fig. 2.5). Note that one endpoint of each obstacle lies on ∂s_{l+1} and the other is one of the endpoints of the edge containing that obstacle. (Assuming general position, no point in Q, and therefore no edge endpoint, lies on the boundary of any region.) Let I_{l+1} be the set of points in $\mathbb{R}^2$ which are contained in the interior of s_{l+1} and which do not lie in

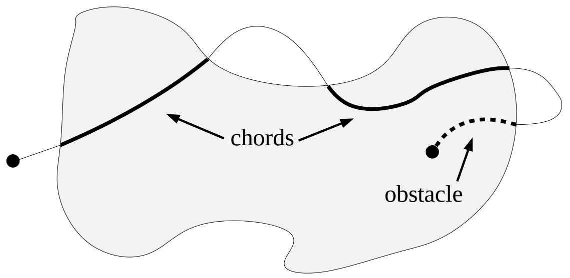

Figure 2.5: Chord and obstacle segments of an edge w.r.t. a region.

the interior of any of the edges in $E_c \cup E'_l$ or the interior of any of the obstacles. Note that no edge in $E_c \cup E'_l$ or an obstacle pierces s_{l+1}. Moreover, since no two obstacles cross, any common point two of them may share will be an edge endpoint belonging to $I_{l+1} \cap Q(s_{l+1})$. Thus, all points in $Q(s_{l+1})$ belong to a connected component I'_{l+1} of I_{l+1}. Therefore, there is a set D' of edges contained in I'_{l+1} such that the plane multigraph $G_{l+1} = (Q, E_c \cup E'_l \cup D')$ properly connects s_{l+1}.

We claim that no edge $e \in D'$ pierces any of $s_1, \cdots, s_l$. For contradiction, assume that some $e \in D'$ pierces some s_i, $i \leq l$. Then, e contains a chord c of s_i that splits s_i into two connected components, each containing points from Q. Since $E_c \cup E_l$ properly connects s_i, there must be an edge $e' \in E_c \cup E_l$, whose endpoints belong to different components of $s_i \setminus c$. This implies that e' crosses c an odd number of times. Since e is an edge in the plane graph G_{l+1}, e' cannot be an edge in G_{l+1}. Therefore, $e' \in E_l \setminus E'_l$, meaning that e' pierces s_{l+1}. All the intersections between e' and c happen at segments of e' which are chords of s_{l+1}, since we excluded the interiors of the obstacles from I_{l+1}. Hence, there is one segment c' of e' which is a chord of s_{l+1} and has an odd number of intersections with c. Moreover, c' lies in the interior of s_i since it is contained in e', and similarly c lies in the interior of s_{l+1}. This contradicts Lemma 4 and therefore e cannot be piercing s_i.

Therefore, none of the edges in $E_c \cup E'_l \cup D'$ pierce any of $s_1, s_2, \cdots, s_l$, and moreover $E_c \cup E'_l \cup D'$ properly connects s_{l+1}. Using the induction hypothesis for regions $s_1, \cdots, s_l$ with $E_c \cup E'_l \cup D$ as the new compulsory set of edges, we obtain a plane multigraph whose edge set properly connects $s_1, s_2, \cdots, s_{l+1}$. □

Lemma 6. *Given a family S of k-admissible regions and a set of points P in the plane, there is a planar graph $G_Q = (Q, E_Q)$ for any $Q \subseteq P$, such that for any $s \in S$, $Q(s)$ induces a connected subgraph.*

Proof. The required planar graph is obtained by applying Lemma 5 for the family S and the set Q, with $E_c = \emptyset$ as the compulsory set of edges and replacing multi-edges with single edges in the resulting plane multigraph. [5] □

Lemma 6 and Theorem 8 imply:

Theorem 15. *The dual range space $\mathcal{R}^*(P, S)$ defined by a family S of k-admissible regions and a set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

The above theorem also follows from [16]. Showing the existence of a $o(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ size ϵ-net for the primal range space $\mathcal{R}(P, S)$ is an open question. However, we prove the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ when S is a family of 2-admissible regions (pseudo-disks). This result was not previously known.

Lemma 7. *Let S be a family of pseudo-disks and P a set of points in $\mathbb{R}^2$. Then, for any $R \subseteq S$ there is a graph $G_R = (R, E_R)$, such that $|E_R| < 24|R|$, and such that for any $p \in P$, if*

[5] An interesting observation is that since planar graphs are 4-colorable [7], this implies that the primal range space defined by a set of points and a set of k-admissible regions in the plane is 4-colorable. The colorability of *dual* range spaces induced by geometric objects has been studied in [48].

p is contained in d pseudo-disks from R, then there are at least $\frac{1}{4}d - 1$ edges among those pseudo-disks.

Proof. For a point $p \in P$, we denote by $R(p) \subseteq R$ the set of pseudo-disks in R that contain p and we define the *degree* $d(p)$ of p as $d(p) = |R(p)|$.

We start by constructing a set $P' \subseteq P$ as follows: Let $p_1, p_2, \cdots, p_n$ be all the points of P in decreasing order of their degrees. Insert p_i into P', if and only if, for every $p_j \in P'$ with $j < i$, we have $|R(p_i) \cap R(p_j)| < \frac{1}{4}d(p_i)$.

Using Lemma 6 for P' and R, we get a planar graph $G' = (P', E')$ such that for each $s \in R$, the points in $P'(s) = \{p \in P' : p \in s\}$ induce a connected subgraph. Therefore, for any region $s \in R$, it holds that $m_s \geq n_s - 1$, where n_s is the number of points of P' contained in s and m_s is the number of edges among those points. Summing over all $s \in R$, we get:

$$\sum_{s \in R} n_s \leq \sum_{s \in R} m_s + |R|. \tag{2.8}$$

Since for any $p_1, p_2 \in P'$, the sets $R(p_1)$ and $R(p_2)$ share fewer than $\frac{1}{4}\min\{d(p_1), d(p_2)\}$ pseudo-disks, we have:

$$\sum_{s \in R} m_s < \sum_{(p_1, p_2) \in E'} \frac{1}{4} \min\{d(p_1), d(p_2)\}.$$

Combining this with the fact that $\sum_{s \in R} n_s = \sum_{p \in P'} d(p)$, Inequality (2.8) gives

$$\begin{aligned}
\sum_{p \in P'} d(p) &< \frac{1}{4} \sum_{(p_1, p_2) \in E'} \min\{d(p_1), d(p_2)\} + |R| \\
&\leq \frac{3}{4} \sum_{p \in P'} d(p) + |R|,
\end{aligned}$$

where for the last inequality we are using the fact that since G' is planar, there is a way to orient the edges in E' so that the out-degree of every node in P' is at most 3 (see [14]). By rearranging, we get:

$$\sum_{p \in P'} d(p) < 4|R|. \tag{2.9}$$

Now we will construct G_R in such a way that for every $p \in P'$, $R(p)$ induces a connected subgraph. Consider some $p \in P'$ and the family $R(p)$ of pseudo-disks that contain it. Using Lemma 2.11 of [2], we obtain a combinatorially equivalent family of pseudo-disks that are all star-shaped with respect to p. From this, by performing an angular sweep using a ray emanating from p, (and having initial orientation such that it doesn't pass through any of the intersection points among the pseudo-disks) we get a combinatorially equivalent family of pseudo-parabolas. Applying Lemma 3 gives a graph $G_p = (R(p), E_p)$, with $|E_p| \leq 6|R(p)| = 6d(p)$, while for any $q \in P$, the pseudo-disks in $R(p) \cap R(q)$ induce a connected subgraph of G_p.

The union of the graphs G_p, for all $p \in P'$ gives the graph $G_R = (R, \cup_{p \in P'} E_p)$. Due to the way that P' was constructed, for any $p \in P \setminus P'$ there is some $p' \in P'$, such that $|R(p') \cap R(p)| \geq \frac{1}{4} d(p)$. Therefore, since the pseudo-disks of $R(p') \cap R(p)$ induce a connected subgraph in $G_{p'}$, there are at least $\frac{1}{4} d(p) - 1$ edges among the pseudo-disks that contain p. On the other hand, if $p \in P'$, then there are at least $d(p) - 1$ edges among the pseudo-disks in $R(p)$. We conclude the proof by observing that the total number of edges in G_R is at most $6(\sum_{p \in P'} d(p)) < 24|R|$ (using (2.9)). $\square$

From Lemma 7 and Theorem 8 we get:

Theorem 16. *The primal range space $\mathcal{R}(P, S)$ defined by a family S of pseudo-disks and a set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Remark 1. *Since pseudo-parabolic halfplanes are a special case of pseudo-disks, Theorem 16 implies that the primal range space defined by a family of pseudo-parabolic half-*

planes and a set of points in the plane, also admits an ϵ-net of size $O(\frac{1}{\epsilon})$.

Remark 2. *Note that the definition of pseudo-disks that we used here is different from the one used in [36] where the family of pseudo-disks is required to be such that there is exactly one pseudo-disk passing any three non-collinear points. We do not make such an assumption and hence our result is stronger than the one in [36]. Moreover, Theorem 16 cannot be proved using the framework developed in [16], since their technique is applicable only to* dual range spaces *(in the sense in which we have used the term in this chapter, see Section 5.2) induced by geometric objects.*

Chapter 3

PTAS for Geometric Hitting Sets and Independent Sets

In Chapter 2, we looked at a technique for proving the existence of ϵ-nets of small size for certain geometric range spaces. Our construction also gave algorithms for computing ϵ-nets of small size. A natural question then is whether we can compute the smallest ϵ-net that a range space admits. In this chapter, we consider a more general problem of computing the smallest hitting set and present polynomial time approximation schemes (PTAS) for range spaces for which only constant factor algorithms (with a rather large constants) were known. The problems we consider are strongly NP-complete and hence, unless P=NP, it is not possible to find a fully polynomial time approximation scheme (FPTAS) for these problems. Quite surprisingly, our algorithm is a very simple local search algorithm which iterates over local improvements only. The proof technique used also yields a PTAS for the maximum independent set in the intersection graph of an r-admissible set of regions in the plane. This extends similar results obtained in [4] and [1]. Finally, the algorithmic technique we use gives a new way to prove the existence of small ϵ-nets for range spaces induced by unit squares in the plane. We believe that a similar proof may exist for more general range spaces.

3.1 Introduction

In the minimum hitting set problem, we are given a range space $\mathcal{R} = (P, \mathcal{D})$ consisting of a set P and a set $\mathcal{D}$ of subsets of P called the *ranges*, and the task is to compute the smallest subset $Y \subseteq P$ which has a non-empty intersection with each of the ranges in $\mathcal{D}$. This problem is equivalent to the set cover problem and is strongly NP-hard. If there are no restrictions on the set system $\mathcal{R}$, then it is known that, unless P=NP, there does not exist any polynomial time algorithm that can approximate the minimum hitting within a factor of $c \log n$ of the optimal [46]. The problem is NP-complete even for the case where each point of P lies in at most two sets of $\mathcal{R}$ [23].

A natural occurrence of the hitting set problem is when the range space $\mathcal{R}$ is derived from geometry. For example, given a set P of n points in $\mathbb{R}^2$, and a set $\mathcal{D}$ of m convex polygons containing points of P, compute the minimum-sized subset of P which hits all the polygons in $\mathcal{D}$. Unfortunately, for most geometric range spaces, computing the minimum-sized hitting set remains NP-hard. For example, even the (relatively) simple case where $\mathcal{D}$ is a set of unit disks in the plane is strongly NP-hard [27]. In this paper, we will only be concerned with set systems where P is a set of points, and the ranges in $\mathcal{D}$ are induced by various geometric objects.

Since there is little hope of computing the minimum-sized hitting set for general geometric problems in polynomial time, effort has turned to approximating the optimal solution. In this regard, an interesting connection to the ϵ-net problem was made by Bronnimann and Goodrich [10].

They proved the following[1]: let $\mathcal{R} = (P, \mathcal{D})$ be a range-space for which we want to compute a minimum hitting set. If we can compute an ϵ-net of size c/ϵ for the weighted ϵ-net problem for $\mathcal{R}$ in polynomial time then we can compute a hitting set of size at most $c \cdot \textsc{opt}$ for $\mathcal{R}$, where $\textsc{opt}$ is the size of the optimal (smallest) hitting set, in polynomial time. A shorter, simpler proof was given by Even *et al.* [22].

This connection between ϵ-nets and computing hitting sets implies that for the ranges mentioned above with $O(1/\epsilon)$-sized nets, there exist polynomial-time, constant-factor approximation algorithms for the corresponding hitting set problems. The constant in the approximation then depends on the constant in the size of the ϵ-nets, which are typically quite large. For example, for $(P, \mathcal{H})$, where $\mathcal{H}$ is the set of halfspaces in $\mathbb{R}^3$, the best constant we get using techniques in Chapter 2 is at least 20, yielding at best a 20-approximation algorithm. Furthermore, this is a fundamental limitation of the technique: it *cannot* give better than constant-factor approximations. The reason is the following: the technique reduces the problem of computing a minimum size hitting set to the problem of computing the minimum size ϵ-net and then uses constant factor approximation for the latter problem. It uses the fact that an ϵ-net of size c/ϵ can always be computed and that $1/\epsilon$ is a lower bound on the size of the ϵ-net to get the constant factor approximation. The Bronnimann-Goodrich technique therefore cannot give a PTAS even for relatively simple hitting set problems.

As a side effect of the techniques we use in this chapter, we also get some improved results for the maximum independent set for geometric intersection graphs. We briefly review the known results in this area. In the maximum independent set problem for geometric

[1] They actually proved a more general statement, but the following is more relevant for our purposes.

intersection graphs we are given a set of geometric objects and the goal is to compute the maximum independent set in the intersection graph defined by them. In other words, we are required to compute a pairwise non-intersecting subset of the objects of the largest size. For general graphs, it is impossible to approximate the maximum independent within a factor better than $n^{1-\delta}$ for any $\delta > 0$, unless NP=ZPP [25]. Even when the graph is an intersection graph of simple geometric objects like unit disks in the plane or orthogonal line segments in the plane, computing the maximum independent set is NP-hard [29]. A PTAS for the unit disks case appeared in [28] following which a PTAS for arbitrary disks appeared in [12] and [21]. These algorithms, however, uses a *shifted dissection* technique which requires the objects to be fat. Agarwal *et. al.* gave a PTAS for the case of unit height rectangles in [4] and more recently, Agarwal and Mustafa [1] gave a polynomial time constant factor approximation algorithm for the case of non-piercing rectangles in the plane.

3.2 Results

We prove the following results in this chapter:

$(1 + \delta)$-approximation of the minimum hitting set via local search. We present a new general technique for approximating geometric hitting sets, which avoids the limitation of the Bronnimann-Goodrich technique: we give the first polynomial-time approximation schemes for the minimum geometric hitting set problem for a wide class of geometric range spaces. All these problems are strongly NP-complete and hence, unless P=NP, there is no FPTAS for these problem. Specifically, we show that:

- Given a set P of n points, and a set $\mathcal{H}$ of m halfspaces in $\mathbb{R}^3$, one can compute a $(1 + \delta)$-approximation to the smallest subset of P that hits all the halfspaces in $\mathcal{H}$ in

$O(mn^{O(\delta^{-2})})$ 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 $(1+\delta)$-approximation to the smallest subset of P that hits all the regions in $\mathcal{D}$ in $O(mn^{O(\delta^{-2})})$ time. This includes pseudo-disks (they are 2-admissible), same-height rectangles, circular disks, translates of convex objects etc. See Definition 3.3.1 for the definition of an r-admissible set of regions.

The above results should be contrasted with the fact that even for relatively simple range spaces like those induced by unit disks in the plane, the previous best known approximation algorithm is due to a recent paper of Carmi *et. al.* [11] which gives a 38-approximation algorithm improving the earlier best known factor of 72 [42].

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 17. *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. Then, there exists a constant c such that a $(c/\delta)^2$-level local search algorithm returns a hitting set of size at most $(1+\delta)\cdot \textsc{opt}$, where $\textsc{opt}$ is the size of the optimal(smallest) hitting set.*

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 $O(mn^{2k+1})$. By using data-structuring techniques, this bound can be improved by polynomial factors; however that is not our goal here.

As a part of proving Theorem 17, we prove a result about planar bipartite graphs, which is of independent interest. 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., $N_G(W) = \bigcup_{v \in W} N_G(v)$. We prove the following:

Theorem 18 (Planar Expansion Theorem)**.** *There exists constants c and k_0 such that for any $k \geq k_0$, if $G = (R, B, E)$ is a bipartite planar graph on red and blue vertex sets R and B, $|R| \geq 2$, so that for every subset $B' \subseteq B$ of size at most k, $|N_G(B')| \geq |B'|$. then* $|B| \leq (1 + c/\sqrt{k})\, |R|$.

$(1 + \delta)$-approximation of maximum independent set in geometric intersection graphs. We give a PTAS for the maximum independent set of the intersection graph defined by a set of r-admissible set of regions in the plane. This extends the results obtained in [4] and [1]. Our algorithm is again a k-level local search similar to the one used for approximating minimum hitting sets. We start with the empty set and repeatedly try to replace $k - 1$ or fewer objects by a larger number of objects so that the resulting set of objects is still pairwise non-intersecting. We again use Theorem 18 to prove the following:

Theorem 19. *Let $\mathcal{D}$ be an r-admissible set of regions in the plane. There is a constant c such that a (c/δ^2)-level local search returns an independent set of the intersection graph of $\mathcal{D}$ of size at least* $\textsc{opt}/(1 + \delta)$ *where* $\textsc{opt}$ *is the size of the maximum independent set.*

Existence of small ϵ-nets via local search. We show that the local search technique can also be used to prove the existence of small size ϵ-nets. Specifically, we show that for the case where we have points in the plane and the 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. So far, the only other place where local search has been used to prove a bound on the size of ϵ-nets is the proof of the existence of $O(1/\epsilon)$ size ϵ-nets for halfspaces in $\mathbb{R}^2$ by Pach and Woeginger [44]. It is not at all clear that the same holds for halfspaces 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.

In Section 3.3 we present the proof of the main Theorem 17, assuming Theorem 18. We prove Theorem 19 in Section 3.4 again assuming Theorem 18. Section 3.5 then gives the proof of the Planar Expansion Theorem (Theorem 18).The alternate proof for the existence $O(1/\epsilon)$ size ϵ-nets for unit squares in the plane is given in Section 3.6.

3.3 PTAS for minimum hitting 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 $\delta > 0$, an $O(\delta^{-2})$-level local search returns a hitting whose size is at most $(1 + \delta)$ 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 $D \in \mathcal{D}$, there is an edge between a vertex in $D \cap R$ and a vertex in $D \cap B$ whenever both the intersections are non-empty.

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.

Let us now return to the general problem. Let $\mathcal{R} = (P, \mathcal{D})$ be a range space satisfying the locality condition where P is a set of size n and $\mathcal{D}$ is a set of m subsets of P. Let $R \subseteq P$ be a hitting set of minimum size and let B be a hitting set returned by a k-level local search. We will use the fact that no local improvement is possible in B to show that $|B|$ cannot be too much larger than $|R|$.

We can 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 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 $|R|$.

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 4. *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 $B \setminus B'$. □

Claim 4 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.

Now Theorem 18 implies that given any parameter δ, a k-level local search with $k = c^2\delta^{-2}$ gives a $(1 + \delta)$-approximation to the minimum hitting set problem for $\mathcal{R}$.

3.3.1 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 [45], 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. Recall the definition of an r-admissible set regions:

Definition 3.3.1. *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 $l \leq k$ 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 (which are 2-admissible), same-height rectangles, circular disks, translates of convex objects etc.

3.3.2 PTAS for halfspaces in $\mathbb{R}^3$.

Given a set of halfspaces 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 halfspaces containing it. Let $\mathcal{R} = (P, \mathcal{D})$ be the range space defined by the remaining set of points and the remaining set of halfspaces. 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 ∂C of the convex hull C of $R \cup B$. In the first stage, we add all red-blue edges (1-faces) of C to G. In the second stage we map each red or blue point p lying in the interior C to a triangular face $\Delta(p)$ of C which intersects the ray op emanating from o and passing through p.[2] 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 it 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

[2]Here we are assuming that each face of C is a triangle, since one can always triangulate the faces.

clear that this can be done while keeping the graph planar. For each blue point $b \in Q$, we try to find one corner c of Δ such that there is no halfspace $h \in \mathbb{R}^3$ with the following properties: i) the only blue point in h is b, ii) h contains exactly one of the corners of Δ. 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 halfspace 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 [35] 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 halfspace 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 halfspace $h \in \mathbb{R}^3$, that does not contain o and that 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 ∂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 ∂C. Consider the halfspace h' parallel to and contained in h that contains the smallest number of points and still contains both red and blue points. Clearly, 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 connected to c since it cannot be the case that we connected b to the other two corners of Δ. Hence, in all cases, b is connected to one of the red points in h'.

3.4 PTAS for maximum independent set

Let $\mathcal{D}$ be an r-admissible set of regions in the plane. We want to approximate the maximum independent set in the intersection graph of these regions. Let R be a maximum independent set and let B be the independent set returned by a k-level local search. The regions in R are pairwise non-intersecting and so are the regions in B. Since $R \cup B$ also forms an r-admissible set, the intersection graph G of the regions in $R \cup B$ is a planar bipartite graph. Since no local improvement is possible in B, i.e. no subset of $B' \subseteq B$ of size $k-1$ can be replaced by a set of size k so that the resulting set if still pairwise non-intersecting, we conclude that for every subset $R' \subseteq R$ of size at most k, $|N_G(R')| \geq |R'|$. Applying Theorem 18 on G with the roles of R and B exchanged, we have: $|B| \geq |R|/(1 + c/\sqrt{k})$ for some constant c. This proves Theorem 19.

3.5 Proof of the Planar Expansion Theorem

We will use the following result for the proof of the theorem:

Theorem 20 (Planar graph partition with small boundary size [33])**.** *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 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 18. Let $r = |R|$ and $b = |B|$. Consider the groups of G formed according to Theorem 20 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 20 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., $b_i^{int} \leq r_i^{int} + r_i^\partial$. Adding b_i^∂ to both sides and summing over all i we have

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

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

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

□

3.6 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 $|Y| = O(1/\epsilon)$.

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 griding 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 5. $|A| = O(1/\epsilon)$.

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. □

Each unit square $s \in S$ contains exactly one of the grid points, and for the squares in

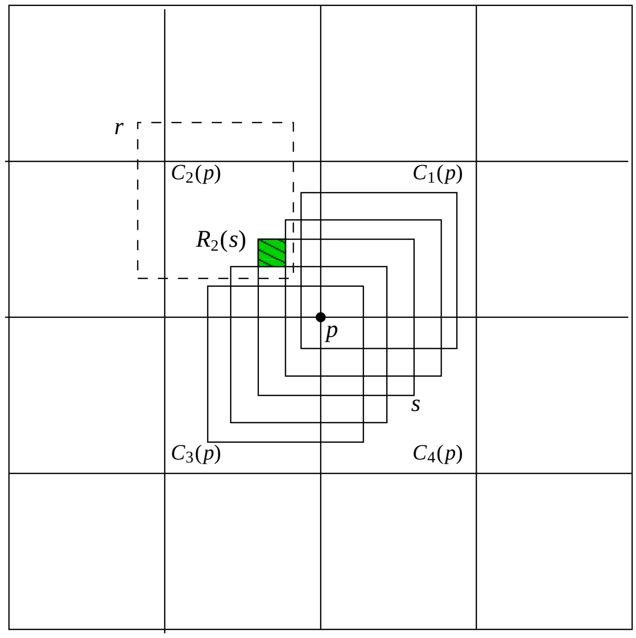

Figure 3.1: The normal square r covers the ϵ-net square s and stabs its neighbors (in the cascade $M_2(p)$) in the cell $C_2(p)$.

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 3.1). 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 8. *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 $C_i'(p)$.*

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. □

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 $Z_i(p) = \cup_{s \in M(p)} N(s)$ be the neighborhood of $M_i(p)$.

Claim 6. *For any $p \in A$ and any i, $|Z_i(p)| \geq \frac{|M_i(p)|}{15} \cdot \epsilon n$. Furthermore, $|M(p)| \leq 60 \frac{|Z(p)|}{\epsilon n}$.*

Proof. First notice that the second statement in the claim easily follows from the first since for some j, $M_j(p) \geq M(p)/4$ and therefore $|Z(p)| \geq |Z_j(p)| \geq |M_j(p)|/15 \geq |M(p)|/60$. We will now prove the first statement.

Partition the squares in $M_i(p)$ into two types: those that have personal dense points in $C_i(p)$ only, or in both $C_i(p)$ and $C_i'(p)$. If the former set has size at least $|M_i(p)|/3$, we are done: each such square has $N(s) \geq \epsilon n$ (due to the personal dense point), and by Lemma 8, each normal square is double-counted at most five times when summing up $N(s)$ for squares in this set. Therefore $|Z_i(p)| \geq (|M_i(p)|/15)\epsilon n$.

Otherwise, assume that there are at least $2|M_i(p)|/3$ squares, say set M', which have personal dense points in both $C_i(p)$ and $C_i'(p)$. Let $t = |M'|$ and let $s_1, s_2, \cdots, s_t$ 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 $C_i(p)$ ($C'_i(p)$). 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 $C_i(p)$ ($C'_i(p)$). Lemma 8 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 fifth 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 $C_i(p)$ since otherwise one of them would have a far red successor. Now, since $|M'''| \geq |M_i(p)|/15$ and each normal square can contain the personal dense point of at most one of the squares of M''' in $C_i(p)$, we have $|Z_i(p)| \geq (|M_i(p)|/15)\epsilon n$.

□

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 second inequality in Claim 6 over all $p \in A$ and using Claim 5, one gets the required statement: $|M| = \sum_{p \in A} |M_p| = O(1/\epsilon)$.

3.7 Future Work

We gave a PTAS for some geometric hitting set problems and proved a theorem, of independent interest, about bipartite planar graphs in the process. We believe that the theorem about bipartite planar graphs may be true for a more general class of graphs. This may allow us to get PTAS for other geometric hitting set problems. It is also worth exploring whether there is a PTAS with a running time $O(mn^{O(\epsilon^{-1})})$ instead of $O(mn^{O(\epsilon^{-2})})$ for the problems we considered.

We also believe that the local search technique can be used to find alternative proofs of the existence of small ϵ-nets for many other geometric range spaces including those induced by half-spaces in $\mathbb{R}^3$ and by an r-admissible set of regions in the plane. Currently, however, it is not even clear how to use it to prove $O(1/\epsilon)$ size ϵ-nets for range spaces induced by different sized squares in the plane.

Chapter 4

Small weak ϵ-Nets

So far we have been discussing strong ϵ-nets where given a range space $\mathcal{R} = (X, R)$, we want to pick a subset of X which *hits* all *ϵ-heavy* ranges. However, as we saw in Chapter 1, if X is a set of points in $\mathbb{R}^d$ and R is the set of all subsets of X obtained by intersecting X with convex sets in $\mathbb{R}^d$, there is no hope of obtaining a strong ϵ-net for $\mathcal{R} = (X, R)$ of size dependent only on ϵ. For example if X is in convex position then any strong ϵ-net N must contain at least $(n - \epsilon n - 1)$ points since otherwise the points not included in N form an ϵ-heavy range which is not hit by N. We then introduced the idea of weak ϵ-nets which allow us to hit ϵ-heavy convex ranges with a small number of points, their number depending only on ϵ. In this chapter we study weak ϵ-nets of constant size. We start with the simplest instance of a weak ϵ-net, namely the centerpoint, whose existence is given by the centerpoint theorem. The centerpoint theorem is one of the fundamental combinatorial results in discrete geometry, with applications in geometric algorithms [15, 41, 50], large-scale computing [40], multivariate data analysis [20] and several others. It states the following:

Centerpoint Theorem [43, 35]. [1] Given a set P of n points in the plane, there exists a point c (not necessarily in P) such that any convex set containing more than $\frac{2}{3}n$ points of P contains c. Furthermore, this bound is tight.

We will look at a generalization of the above theorem to more than one point. For example, is it possible to find two points c_1 and c_2 in the plane such that any $\frac{1}{2}$-heavy (i.e. containing at least $n/2$ points of P) convex set must contain either c_1 or c_2? We present a general procedure that gives the following results: one can hit all $\frac{4}{7}$-heavy convex sets with 2 points. Furthermore, we prove that this bound is tight. Similar results are derived for larger number of points. In particular, we show that 5 points suffice to hit all $\frac{20}{41}$-heavy convex sets. This improves a natural way of picking five points [8] which hit all $\frac{1}{2}$-heavy convex sets: find two lines (using the ham-sandwich theorem [35]) which partition the point set into four regions with $n/4$ points in each. Add the intersection point x of the lines along with the centerpoints of the points in each of the four regions. Since any set avoiding the four centerpoints (one for each region) can contain only $\frac{2}{3}rd$ of the points in any of the regions and must avoid one of the regions completely if it avoids x, these 5 points form a $\frac{1}{2}$-net.

Related work

Aronov *et al.* [8] proved that given a set P of n points in the plane, all convex sets containing more than $\frac{5}{8}n$ points of P can be hit by two points. They also construct inputs where

[1] This theorem can be equivalently stated as: there exists a point c such that any halfspace containing c contains at least $n/3$ points of P.

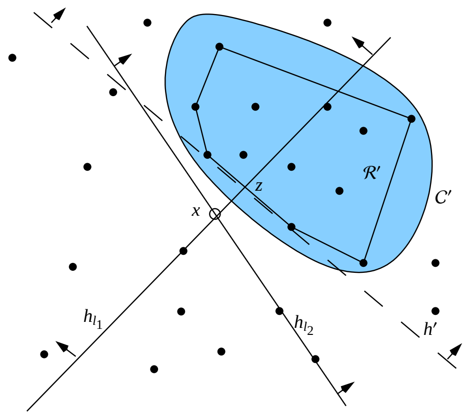

Figure 4.1: Illustration of Theorem 21

regardless of how one picks the two points, there exists a convex set containing at least $\frac{5}{9}n$ points that is not hit. we improve both their results to get the optimal result of $\frac{4}{7}n$. We similarly improve their results for other small numbers of points (see Section 4.2 for specific improvements).

Another related area of research is the so-called *Gallai-type* problems [35] which ask whether certain families of geometric shapes can be "pierced" by a small number of points. An example of such a problem is the following: Given a set of closed disks in the plane such that every pair intersects, what is the smallest number of points needed to hit all these disks? In this case, the answer which is both necessary and sufficient, is four [17]. In our problem, we are looking to hit considerably more general objects(convex sets), with the added constraint that one first fixes n input points, and each convex set contains a constant proportion of these points.

4.1 Main Theorem

We first present some definitions. Given a set P of n points in $\mathbb{R}^d$ and a finite set $Q \subset \mathbb{R}^d$, define the following:

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

and let $\epsilon_i^d(P) = \min_{Q,|Q|=i} \epsilon(P, Q)$. Set $\epsilon_i^d = \sup_P \epsilon_i^d(P)$. In other words, given any P, the set of all convex sets containing $\epsilon_i^d n$ points of P can be hit by i points. These i points are said to form a *weak* ϵ_i^d-net for P. The centerpoint theorem in d dimensions states that $\epsilon_1^d = \frac{d}{d+1}$.

We fix a direction $\overrightarrow{u} \in \mathbb{R}^d$ 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 C, let $f_u(C)$ denote the height of the lowest point in C, i.e. $f_u(C) = \inf_{p\in C} f_u(p)$.

We now present our main result.

Theorem 21. *We have $\epsilon_0^d = 1$ and for $r, s \geq 0$ and $d \geq 1$, we have*

$$\epsilon_{r+ds+1}^d \leq \frac{\epsilon_r^d \cdot (1 + (d-1)\epsilon_s^d)}{1 + \epsilon_r^d \cdot (1 + (d-1)\epsilon_s^d)}.$$

We have $\epsilon_0^d = 1$ since we don't need any point to hit 1-heavy convex sets as there are no such sets. [2] Let a,$b \in [0, 1]$ be two reals to be fixed later. Let $\mathcal{H}$ be the set of all closed halfspaces that contain at least an points of P and whose bounding hyperplane passes through d points in P. Define $\mathcal{H}_b^d = \{(h_{i_1}, h_{i_2}, \cdots, h_{i_d}) \in \mathcal{H}^d \mid |P \cap (h_{i_1} \cap h_{i_2} \cdots \cap h_{i_d})| \geq bn, \}$

[2]Recall that an ϵ-heavy range contains *more* than an ϵ fraction of the elements of the ground set.

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 $(h_{l_1}, \cdots, h_{l_d})$, such that

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

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

We choose the upward direction $\overrightarrow{u}$ so that the d-tuple $(h_{l_1}, \cdots, h_{l_d})$ is well defined. Note that $f_u(h_{i_1} \cap \cdots \cap h_{i_d}) = -\infty$ iff $h_{i_1} \cap \cdots \cap h_{i_d}$ is unbounded in the *downward* direction $-\overrightarrow{u}$. Let $\mathcal{P}$ be the convex hull of P and let $h_{j_1}, \cdots, h_{j_d}$ be d halfspaces defining a vertex v of $\mathcal{P}$ and containing P. Choose the upward direction $\overrightarrow{u}$ so that the vertex v is the unique lowest vertex of the polyhedron $\mathcal{P}' = h_{j_1} \cap \cdots \cap h_{j_d}$ in the upward direction and each of the points $p \in P$ get a unique height. Such a choice of $\overrightarrow{u}$ 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 $-\overrightarrow{u}$. Therefore, $(h_{l_1}, \cdots, h_{l_d})$ is well defined and the lowest point in $h_{l_1} \cap \cdots \cap h_{l_d}$ is unique.

Let $\mathcal{R}$ be the polyhedron $\mathcal{R} = \{h_{l_1} \cap \cdots \cap h_{l_d}\}$. 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 $\{h_{l_1}, \cdots, h_{l_d}\}$ 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 $Q = \{x\} \cup Q' \cup Q_{l_1} \cup \cdots \cup Q_{l_d}$, where

1. x is the unique lowest point in $h_{l_1} \cap \cdots \cap h_{l_d}$.

2. Q' is an ϵ_r^d-net for the point set $P \setminus (P \cap h_{l_1} \cap \cdots \cap h_{l_d})$ using r points.

3. Q_{l_i} is an ϵ_s^d-net for the point set $P \setminus (P \cap R_{l_i})$ using s points.

Lemma 9. *Q is an a-net for P, and has size $r + ds + 1$.*

Proof. The size of Q is obvious from the construction. 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. *Let C' be a convex set containing at least an points of P that does not contain x and contains points from $P \cap h_{l_1} \cap \cdots \cap h_{l_d}$. Then, $|P \cap C' \cap R_{l_i}| < bn$ for some $i \in [1, d]$.*

Proof. For contradiction, assume that C' intersects all R_{l_i} in 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 $h_{l_1} \cap \cdots \cap h_{l_d}$, i) h' intersects $h_{l_1} \cap \cdots \cap h_{l_d}$, 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 $\{h'\} \cup \{h_{l_1}, \cdots, h_{l_d}\} \setminus h_{l_i}$ 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 $(h_{l_1}, \cdots, h_{l_d})$ has the highest lowest-intersection point among the d-tuples in $\mathcal{H}^d$. See Figure 4.1 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 $an \geq \epsilon_r^d(1-b)n$, then C' is hit by Q'.

3. C' does not contain x and yet contains points from $\mathcal{R}$. Then, by Claim 7, $C' \cap R_{l_i} \leq bn$ for some $i \in [1, d]$. Then it must contain at least $an - bn$ points from $P \setminus (P \cap R_{l_i})$. If $an - bn \geq \epsilon_s^d(1 - ((d-1)a - (d-2)))n$, then C' is hit by Q_{l_i}.

Therefore, if

$$an \geq \epsilon_r^d(1-b)n \quad \textit{and} \quad an - bn \geq \epsilon_s^d(d-1)(1-a)n \tag{4.1}$$

then C' is hit by Q. Maximizing a while satisfying (4.1) yields

$$\epsilon_{r+ds+1}^d \leq a = \frac{\epsilon_r^d \cdot (1 + (d-1)\epsilon_s^d)}{1 + \epsilon_r^d \cdot (1 + (d-1)\epsilon_s^d)},$$

completing the proof of Lemma 9 and hence Theorem 21. □

4.1.1 New Proofs of the Centerpoint theorem and Helly's theorem.

The above method actually gives elementary proofs of the centerpoint theorem and Helly's theorem in any dimension.

The proof of the centerpoint theorem in two dimensions is as follows: given a set P of n points, consider all closed halfplanes containing more than $\frac{2}{3}n$ points of P, 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), 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 21 by setting $r = s = 0$ and $d = 2$ to get $\epsilon_1^2 = \frac{2}{3}$. 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 upward direction). The lowest point of their intersection is the centerpoint.

The same idea also gives an elementary proof of Helly's theorem in any dimension. Helly's theorem (see [35]) states that if we have a set of closed convex sets in $\mathbb{R}^d$ and we know that every $d+1$ of them have a common intersection then all of them have a common intersection. To prove this, consider the point p that is the highest lowest-intersection point of any d of the convex sets. By choosing the upward direction carefully, it can be assured that p is uniquely defined and its height is finite. Then, by the condition of the theorem, each of the other convex sets must intersect the common intersection $\mathcal{I}$ of the d convex sets defining p. Hence, they must also contain p since if one of the convex sets C does not contain p then the lowest point q of $C \cap \mathcal{I}$ is the lowest point of the intersection of d of the convex sets and is higher than p. This contradicts our assumption that p is the highest lowest-intersection point of any d of the given convex sets. Therefore, p is lies in all the convex sets and thus they have a common intersection.

4.2 Consequences of main theorem

Improving upon previous work [8], we completely resolve the 2-point case in the plane.

Proposition 2. *Given a set P of n points in $\mathbb{R}^2$, the set of all convex sets which contain more than $\frac{4}{7}n$ points of P can be hit by* two *points (i.e., $\epsilon_2^2 \leq \frac{4}{7}$). Furthermore, there exist arbitrarily large point sets such that the set of all convex sets containing $\frac{4}{7}n$ points cannot be hit by two points.*

Proof. The upper bound follows from Theorem 21 by setting $r = 1$, $s = 0$ and $d = 2$.

Our lower bound construction is similar to the lower bound construction in [8]. We construct a set of P of 7 points 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 4 of the points in P. By replacing each of the points of P by a set of $n/7$ points (for arbitrary n) contained in a sufficiently small disk, one gets a set Q of size n such that no two points in the plane hit all the convex sets containing at least $\frac{4}{7}n$ points of Q.

Our set P is the set of vertices of regular heptagon. Let us name the vertices a, b, c, d, e, f and g in clockwise order. If either p or q is identical to one of the 7 points, say a, then the other point cannot hit the convex sets $bcde$, $defg$ and $fgbc$ simultaneously since they don't have a common intersection. On the other hand, if neither p nor q is identical to any of the 7 points, then one of the closed halfspaces defined by the line passing through p and q contains 4 of the points of P whose convex hull is not hit by either p or q. □

Proposition 3. *Given P, the set of all convex sets which contain more than $\frac{8}{15}n$ points of P can be hit by* three *points (i.e., $\epsilon_3^2 \leq \frac{8}{15}$). Furthermore, there exist arbitrarily large point sets such that the set of all convex sets containing $\frac{5}{11}n$ points cannot be hit by three points.*

Proof. The upper bound follows from Theorem 21 by setting $r = 2$, $s = 0$ and $d = 2$.
The lower bound construction is as follows. We construct a set of 11 points such that for any three given points p, q and t in the plane there is a convex set containing 5 points from P and avoids all the three points. As in the proof of Proposition 2, one can replace each of these points with a set of $n/11$ points (for arbitrary n) contained in a sufficiently small disk and obtain a set Q of points such that no three points in the plane hits all the convex sets containing at least $\frac{5}{11}n$

Our set P is shown in Figure 4.2(a). Assume that there are three points which hit all convex sets containing $\frac{5}{11}n$ points of P. We first show that none of these points can be identical to any of the 11 sets in the point set. Observe that if all the three points are identical to one of the 11 sets in the point set, then they cannot hit the convex hull of the remaining points, of which there are at least 8. Also, if two of the points p, q and t are identical to one of the points, then the remaining points, of which there are at least 9, can be used to define two convex sets containing 5 points each and sharing only one of the 11 points. A single point hitting both these sets should be identical to the shared point implying that all the three points are identical to one of the points. If only one of the points, say p, is identical to one of the 11 points, say the point k, then consider the convex sets $defgh$, $fghij$ and $jabcd$. Since q and t hit all the three sets, one of the points should be contained in the region $hvfg$, where v is intersection point of the segments fj and dh. Now, consider the sets $hijab$ and $bcdef$. The third point must hit both these sets and therefore must be identical to b.

Assuming that none of the points is identical to one of the 11 points, we show that if there exists a set of three points which hits all convex sets containing 5 points from P then one of those points is contained in one of the bold triangles shown in Figure 4.2(a). Consider the four convex sets $jkabc$, $abcde$, $defgh$ and $ghijk$ (see Figure 4.2(b)) containing 5 points each. In order to hit all the four sets, one of the three points must be in one the four regions jzk, gxh, dve or $abcs$. 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 region $abcs$. 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 $abcs$ is in one of the triangles atb or buc (see Figure 4.2(c)), we are done again. So, we assume that it is in the region $stbu$ but does not lie on bt or bu. 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 4.2(c)) since we have assumed that none of the points is identical to f.

Hence, one of the bold triangles shown in Figure 4.2(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 4.3(a)). Now, since the regions $abcjk$ and $abcde$ must be hit by one point (see Figure 4.2(a)), the region $abcs$ contains a point.

Also, since the regions $abijk$ and $bcdef$ must be hit (see Figure 4.3(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 4.3(c)) although it contains 5 points of P. Hence, it is not possible to hit all the convex regions containing 5 points of P using 3 points. □

Aronov et al. [8] proved that $\epsilon_4^2 \leq \frac{4}{7}$. We actually are able to hit sets containing $\frac{4}{7}n$ points by just two points (Proposition 2). Theorem 21 yields $\epsilon_4^2 \leq \frac{16}{31}$, again improving upon Aronov et al.'s result. Improving upon a result of Alon et al. [5], Aronov et al. [8] showed that if each convex set contains $n/2$ points, then they can be hit by five points. Theorem 21 yields an improvement (set $r = 2$, $s = 1$, and $d = 2$).

Corollary 1. $\epsilon_4^2 \leq \frac{16}{31}$.

Corollary 2. $\epsilon_5^2 \leq \frac{20}{41}$.

4.3 Conclusions

We presented a general technique for constructing small number of points that hit all convex sets containing certain fractions of points of P. This then gives an optimal extension of the centerpoint to two points and improves the previous bounds for larger number of points. One intriguing open problem is whether the bound can be closed for the three-point case. Our work leaves a gap ($\frac{5}{11} \leq \epsilon_3^2 \leq \frac{8}{15}$), and it would be nice to get an optimal bound there.

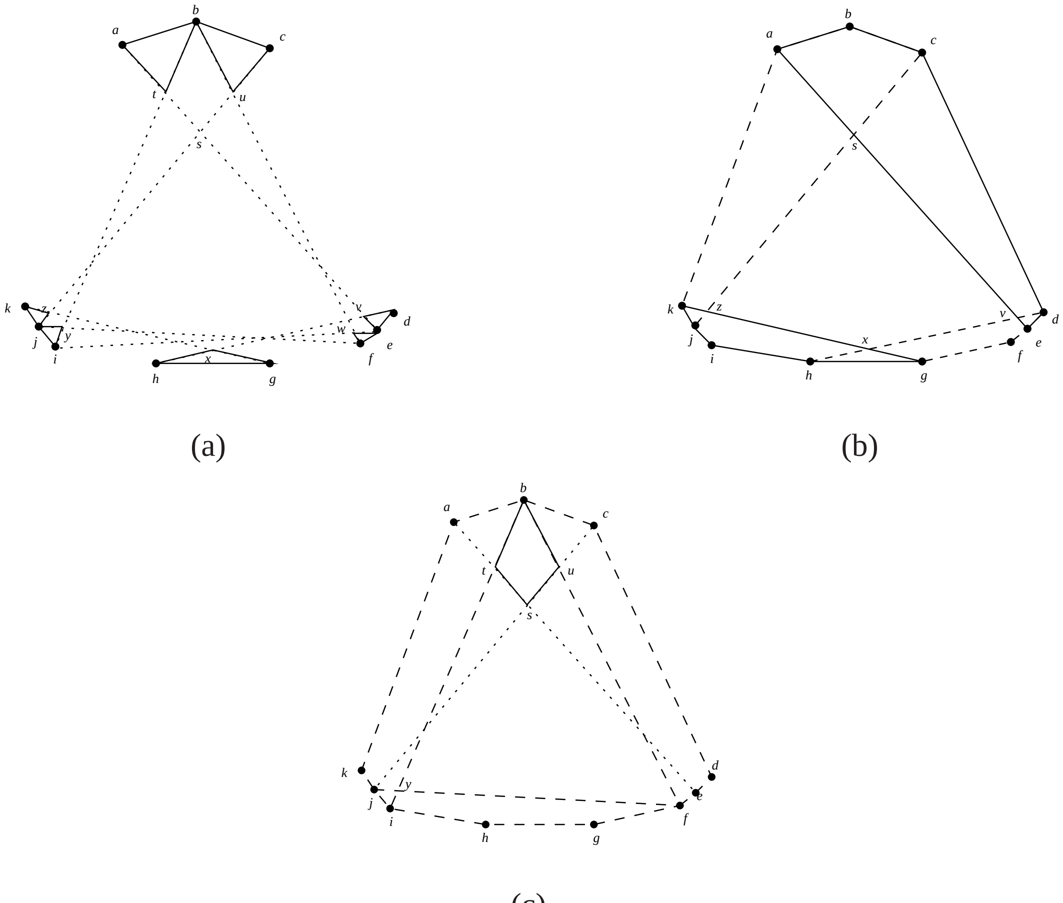

Figure 4.2: (a) One of the seven (bold) triangles contains a point of the weak ϵ-net (b) One of the four regions *jzk*, *gxh*, *dve* or *abcs* contains a point of the weak ϵ-net (c) *jyi* contains a point of the weak ϵ-net

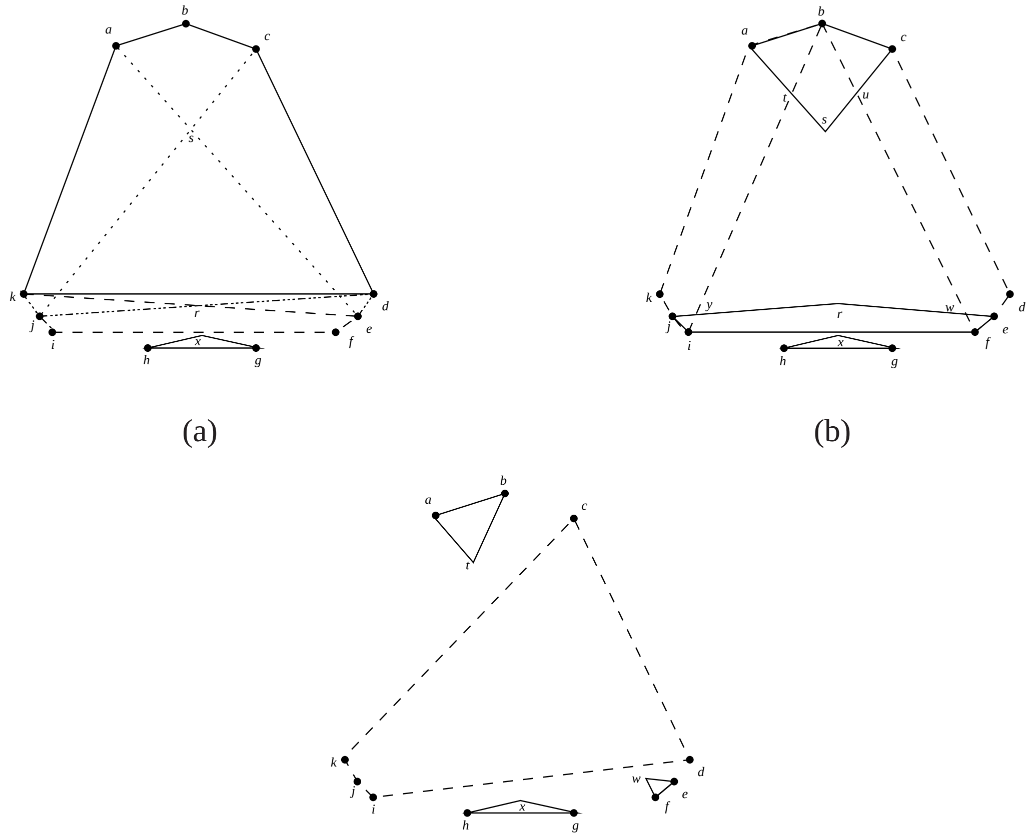

(a) (b)

(c)

Figure 4.3: (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.

Chapter 5

Small basis for weak ϵ-Nets

5.1 Introduction

Given a set system $(X, \mathcal{R})$, where X is the base set, and $\mathcal{R}$ is a family of subsets of X, the strong ϵ-net problem asks for a small subset X' of X such that for every set $S \in \mathcal{R}$ containing at least $\epsilon|X|$ elements, $X' \cap S \neq \emptyset$. As we saw in Chapter 1, if the set system has a finite VC dimension d, then picking a random sample from X of size $O(d\epsilon^{-1} \log(d\epsilon^{-1}))$ yields an ϵ-net with some constant positive probability. We also saw that such small strong ϵ-nets do not exist for set systems of infinite VC dimension. In particular, it fails for range space induces by a set of points and a set of convex objects in the plane.

In this chapter we will be concerned only about weak ϵ-nets with respect to convex ranges. Recall from Chapter 1 that given a finite set P of points in $\mathbb{R}^d$, a weak ϵ-net with respect to convex ranges is a set W of points in $\mathbb{R}^d$, not necessarily belonging to P, such that W has a non-empty intersection with every convex range in $\mathbb{R}^d$ containing more than an ϵ fraction of the points of P. We briefly review a few other basic things that we discussed in

Chapter 1. Let $w(d, \epsilon)$ denote the maximum size of the weak ϵ-net required for any set of points in $\mathbb{R}^d$ with respect to convex ranges. This is finite since Alon *et al.* [5] have shown that for any ϵ, d, there exist a weak ϵ-net of size independent of n. In particular, they proved that $w(d, \epsilon) \leq O(\epsilon^{-(d+1-\delta_d)})$, where δ_d tends to zero with $d \to \infty$. This result was improved by Chazelle *et al.* [13] to $w(d, \epsilon) \leq O(\epsilon^{-d}\text{polylog}(\epsilon^{-1}))$. They also showed that for a set of points in $\mathbb{R}^2$ in convex position, there exists a weak ϵ-net of size $O(\epsilon^{-1}\ \text{polylog}(\epsilon^{-1}))$. More recently, Matousek and Wagner [37] gave an elegant algorithm that computes weak ϵ-nets in $\mathbb{R}^d$ of size $O(\epsilon^{-d}\text{polylog}(\epsilon^{-1}))$.

A long-standing open problem has been to show the existence of weak ϵ-nets in $\mathbb{R}^d$ with size $o(\epsilon^{-d})$. Note that this contrasts sharply with ϵ-nets for finite VC-dimension ranges, where the size of the ϵ-net depends *almost linearly* on $1/\epsilon$. In fact, the current conjecture by Matousek *et al.* [37] is that optimal weak ϵ-nets should have size $O(\epsilon^{-1}\ \text{polylog}(\epsilon^{-1}))$ in $\mathbb{R}^d$ for every integer d. This conjecture and the following observation (which follows from Lemma 15) is the motivation for our work:

Observation 1. *Given a set P of n points in $\mathbb{R}^d$, a weak ϵ-net of P of size k is completely described by $O(d^2k)$ points of P.*

Essentially, each point of the weak ϵ-net is locally constructed from $O(d^2)$ points of P. Hence if weak ϵ-nets do have size $O(\epsilon^{-1}\text{polylog}(\epsilon^{-1}))$ in any dimension, then there must exist $O(\epsilon^{-1}\text{polylog}(\epsilon^{-1}))$ (hidden constants depend on d) points of P from which it is constructed (we call this set a *basis*). So a possible first step towards confirming the conjecture is to show this linear dependence on points of P. *Unfortunately all known constructions of weak ϵ-nets use $\Omega(\epsilon^{-d})$ input points.* In fact, a modification of [37] to compute the weak ϵ-net at one step (instead of several recursive steps) seemed to use fewer input points. However, it does not. Briefly, the construction uses an r-simplicial partition with sets of size $\Theta(n/r)$ such that no hyperplane intersects more than $O(r^{1-1/d})$ sets of the partition. From each set in the partition, one point is chosen and then a set of points, containing a centerpoint for every subset of the chosen r points, is computed. It is then shown that if a convex set intersects $\Omega((d+1)r^{1-1/d})$ sets in the partition then one of the centerpoints computed is contained in the set, for otherwise there exists a hyperplane intersecting $\Omega(r^{1-1/d})$ sets. The case in which the convex set intersects fewer than $O(d+1)r^{1-1/d})$ is dealt with recursively. To avoid recursion, we must choose r in such a manner that $O((d+1)r^{1-1/d})$ sets contain fewer that ϵn points. Since the sets are of size $\Theta(n/r)$, we require that $(d+1)r^{1-1/d}n/r < \epsilon n$ implying that $r > ((d+1)/\epsilon)^d$. Hence, in that case too $\Omega(\epsilon^{-d})$ input points are used.

The main features of the results presented in this chapter are as follows:

- We answer the above question in the affirmative, showing that for every point set P, there exists a set of $O(\epsilon^{-1}\log(\epsilon^{-1}))$ points in $\mathbb{R}^d$ from which one can construct a weak ϵ-net for P. So, while the size of weak ϵ-nets that we compute is $\Theta((\epsilon^{-1}\log(\epsilon^{-1}))^{d^2})$, their description (i.e., points used to construct them) is in fact near-linear in $1/\epsilon$.

- The proof establishes an interesting relation between strong ϵ-nets and weak ϵ-nets. Random sampling works for strong ϵ-nets since the number of ranges is polynomially bounded, and seems doomed when the ranges are exponential in number (since then one requires the probability of not hitting a range to be exponentially small as well). We show that sampling approaches work *if* one takes some 'products' over the sampled points. In particular, we show the following. In $\mathbb{R}^2$, take an ϵ-net with respect to the intersection of every six halfplanes. Then *only* from these $O(\epsilon^{-1}\log(\epsilon^{-1}))$ points, one can construct a weak ϵ-net of size $O(\epsilon^{-3}\log^3(\epsilon^{-1}))$. Similarly, we show that by random sampling $O(\epsilon^{-1}\log(\epsilon^{-1}))$ points in $\mathbb{R}^3$, and taking some function of them, one gets a weak ϵ-net of size $O(\epsilon^{-5}\log^5(\epsilon^{-1}))$. For P in $\mathbb{R}^d$, take a random sample of size $O(\epsilon^{-1}\log(\epsilon^{-1}))$ (with only the constant depending on d). Then another product function of these sampled points yields an ϵ-net with size $O(\epsilon^{-d^2})$.

- Our approach directly relates the size of the weak ϵ-nets to the 'description complexity' of these 'product' functions. We use two 'product' functions over points of P: Radon points, and centerpoints. Our proof reveals the following connection (see Corollary 3 for a stronger statement): let Q be a set of m points in $\mathbb{R}^d$, and let $c(Q)$ be a set of points such that a centerpoint of every non-empty subset of Q is present in $c(Q)$. Then if $c(Q)$ has size $O(m^t)$, one can construct weak ϵ-nets of size $O(\epsilon^{-t}\log^t(\epsilon^{-1}))$. Therefore, showing $t < d$ will lead to an improvement in the size of weak ϵ-net obtained.

We first present an elementary proof for the two-dimensional case in Section 5.3. 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 5.4. The general construction for arbitrary dimensions is then presented in Section 5.5.

5.2 Preliminaries

We review a few concepts from discrete geometry for later use [35].

VC-dimension and ϵ-nets [35] Given a range space $(X, \mathcal{R})$, a set $X' \subseteq X$ is *shattered* if every subset of X' can be obtained by intersecting X' with a member of the family $\mathcal{R}$. The VC-dimension of $(X, \mathcal{R})$ is the size of the largest set that can be shattered. The ϵ-net theorem (Welzl and Haussler [26]) states that there exists an ϵ-net of size $O(d\epsilon^{-1}\log(\epsilon^{-1}))$ for any range space with VC-dimension d.

Radon's theorem [35] Any set of $d+2$ points in $\mathbb{R}^d$ can be partitioned into two sets A and B such that $conv(A) \cap conv(B) \neq \emptyset$.

Ramsey's theorem for hypergraphs [19] There exists a constant $R(n)$ such that given any 2-coloring of the edges of a complete k-uniform hypergraph on at least $R(n)$ vertices, there exists a subset of size n such that all edges induced by this subset are monochromatic.

Hadwiger-Debrunner (*p*,*q*)-theorem [6] Given a set S of convex sets in $\mathbb{R}^d$ such that out of every $p \geq d+1$ set, there is a point common to $q \geq d+1$ of them, then S has a hitting set of finite size and the minimum size of such a set is denoted by $HD_d(p,q)$ (independent of $|S|$).

5.3 Two Dimensions

Consider the range space $\mathcal{R}_k = (P, R)$, where P is a set of n points in the plane, and $R = \{P \cap \bigcap_{i=1}^{k} h_i, h_i \textit{ is any halfspace}\}$ 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 the result of Haussler and Welzl [26], it follows that a random sample of size $O(\epsilon^{-1}\log(\epsilon^{-1}))$ is an ϵ-net for $\mathcal{R}_k$ with some constant probability. Let Q be such an ϵ-net. We have the following structural claim which establishes a relation between strong ϵ-nets and weak ϵ-nets.

Lemma 10. *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 set C in the plane containing at least ϵn points of P, either a) $C \cap Q \neq \emptyset$, or b) there exist $\lfloor k/2 \rfloor$ points of Q in convex position, say $q_i \in Q, i = 1 \ldots \lfloor k/2 \rfloor$, such that C intersects the edge $\overline{q_i q_j}$ for all $1 \le i < j \le \lfloor k/2 \rfloor$.*

Proof. Assume $C \cap Q = \emptyset$. We then give a deterministic procedure that always finds $\lfloor k/2 \rfloor$ such points. W.l.o.g. assume that the convex set is polygonal (since there is always a polygonal convex set $C' \subseteq C$ such that $C' \cap P = C \cap P$), and denote its vertices in cyclic order by $p_1, \ldots, p_m$ 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 with apex at p_i, and extending through p_{i+1} to infinity (define $\overrightarrow{p_{i+1} p_i}$ likewise). See Figure 5.1 (a). Let $T(i, j)$ be the region bounded by $\overrightarrow{p_{i-1} p_i}$, the segments $p_i p_{i+1}, \ldots, p_{j-1} p_j$, 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 $T(i_l, j)$ contains a point of Q.

This process ends when $j = 1$. Assume we have l points $q_1, \ldots, q_l$, together with the indices $i_1, \ldots, i_l$. 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 5.1(b).

Claim 8. *The region $T(i_{t-1}, i_t - 1)$ contains 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. □

Define h_t to be the halfspace incident to the edge $p_{i_t-1}p_{i_t}$ and containing C. Claim 8 immediately implies the following.

Claim 9. *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 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 $t = 1 \ldots, l$. By Claim 9, each halfspace h_t separates all the points in $T(i_{t-1}, i_t)$ from C. Thus all the points of Q are now separated by this intersection (see Figure 5.1 (a) for the separating halfplanes), and since each halfspace contains C, the intersection contains at least ϵn points of P. This contradicts the fact that Q was an ϵ-net to the range space $\mathcal{R}_k$.

Finally, note that the sequence q_t of points obtained, $t = 1 \ldots k$, has the property that the intersection point of any (properly intersecting) pair of segments joining non-consecutive points, lies inside 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 C. Therefore picking every alternate point yields the desired set. □

Set $k = 8$, and compute the ϵ-net for the range space $\mathcal{R}_8$. It follows from Lemma 10 that if a convex set C is not hit by the computed ϵ-net,then there exists a sequence of four points, say a, b, c, d, such that 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: the weak ϵ-net consists of an ϵ-net, say Q, for $\mathcal{R}_8$, and the intersection points of all segments between pairs of points of Q. By the above argument, each convex set containing at least ϵn points of P either contains a point from Q or one of the intersection points. The number of points in the weak ϵ-net constructed above are $O(\epsilon^{-4}\log^4(\epsilon^{-1}))$. We now show that by a more careful argument, this can be reduced to $O(\epsilon^{-3}\log^3(\epsilon^{-1}))$.

Theorem 22. *Given a set P of n points in the plane, construct an ϵ-net Q for the range space $\mathcal{R}_{12}$. 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 $\angle abc$ with the line segment ac to Q'. Then $Q' \cup Q$ has size $O(\epsilon^{-3}\log^3(\epsilon^{-1}))$ and forms a weak ϵ-net for P.*

Proof. Fix a convex set C containing at least ϵn points of P. We may assume that C does not contain any point of Q. Then, from Lemma 10, 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 alternating points, say $\triangle ace$ and $\triangle bdf$. The sum of the interior angles of the two triangles is 2π. By the 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, $\angle fab$.

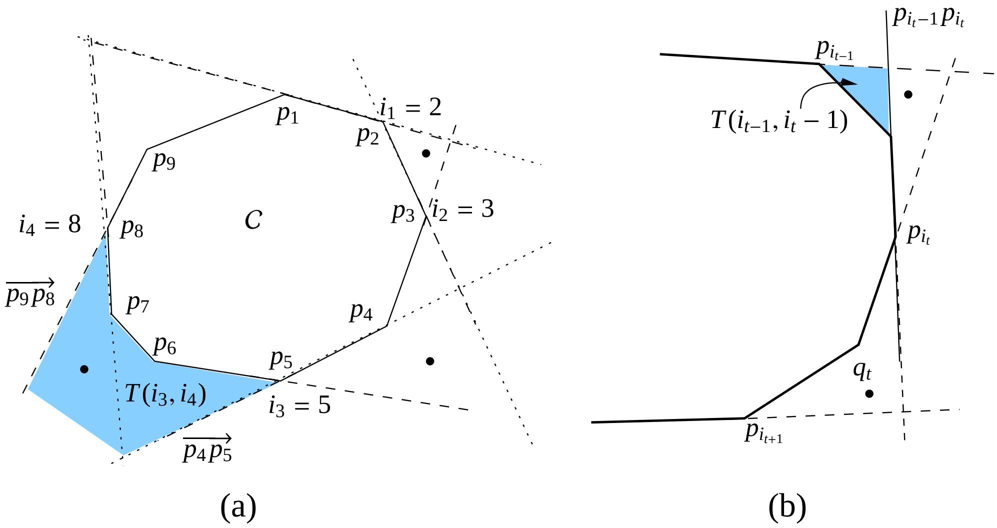

Figure 5.1: Constructing weak ϵ-nets in two dimensions. (a) The dotted lines indicate the at most k halfspaces that are used to separate Q from C.

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 5.2(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'. □

5.4 Three Dimensions

Lemma 11. *For every d and $t \geq d+1$, there exists a constant $f_d(t)$ such that given a polytope C and a set of points Q in $\mathbb{R}^d$ such that $C \cap Q = \emptyset$, i) either there is set of $f_d(t)$ hyperplanes such that each $q \in Q$ is separated from C by one of the 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. Assume, without loss of generality, that the origin lies in the interior of C. For $q \in Q$

define

$$S(q) = \{a \in \mathbb{R}^d | \langle a, q\rangle \geq 1, \langle a, x\rangle \leq 1 \ \forall x \in C\},$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product. First note that $S(q) \neq \emptyset$ since $q \notin C$. Second, $S(q)$ is convex and closed, as it is the intersection of a family of closed convex sets (namely the closed halfspaces defined by the dual of q and the duals of the vertices of C). Since C contains the origin, $S(q)$ is also bounded and hence compact.

Since $0 \notin S(q)$, $a \in S(q)$ implies that there is a hyperplane ($\langle a, x\rangle = 1$) which separates the point q from the C. If there are $d+1$ points $q_1, \cdots, q_{d+1}$ whose convex hull does not intersect C, then these $d+1$ points can be separated from C by a single hyperplane (separation theorem, [35]). This implies that the corresponding convex sets $S(q_1), \cdots, S(q_{d+1})$ have a common intersection.

Let $S = \{S(q) \mid q \in Q\}$ be the set of convex sets 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 C, then $d+1$ of every t convex sets 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 C using $f_d(t)$ hyperplanes, where $f_d(t) = HD_d(t, d+1)$ and $HD_d(p, q)$ is the Hadwiger-Debrunner hitting set number for p and q in d dimensions. □

Lemma 12. *For every $t \geq 5$, there exists a constant $g(t)$ such that given a convex set C in $\mathbb{R}^3$ and set Q' of $g(t)$ points in $\mathbb{R}^3$ so that the convex hull of every 4 of the points in Q' intersects C, one can find $Q'' \subseteq Q'$ of size at least t such that the convex hull of every 3 of the points in Q'' intersects 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 C and 'blue' otherwise. Then, by Ramsey's theorem for hypergraphs [19], 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 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 having disjoint interiors. Since both these tetrahedra must intersect C, a triangular face of each of them must also intersect C (not necessarily in the interior), 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 C. □

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

Claim 10. *Let $T = \{a, b, c, d, e\}$ be a set of five points in convex position in $\mathbb{R}^3$. Then, if a convex set 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 C must intersect ab. Otherwise, there exists a hyperplane h separating ab from C. Since ab intersects the convex hull of c,d and e, h separates at least one point in $\{c, d, e\}$ from C and convex hull of a,b and this third point does not intersect C, a contradiction. □

Lemma 13. *For every t, there exists an $h(t)$ such that given a convex set C in $\mathbb{R}^3$ and a set Q'' of $h(t)$ points so that the convex hull of every 3 points in Q'' intersects C, one can find a subset $Q''' \subseteq Q''$ of size t such that the convex hull of every two points in Q''' intersects C.*

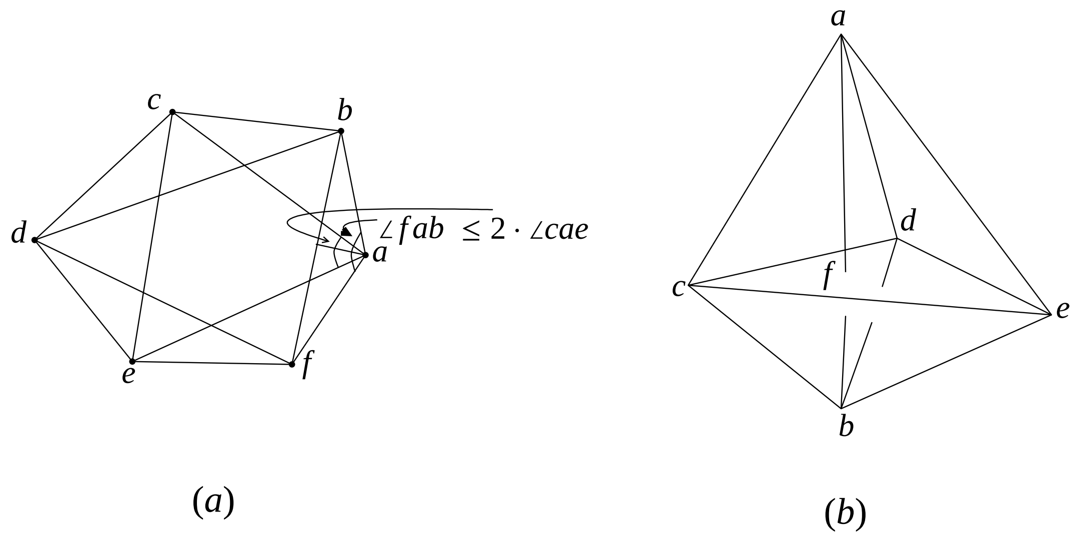

Figure 5.2: (a) The intersection of a bisector with a segment will lie inside C, (b) If C intersects edges ac, ad and ae, then it must intersect af. Similarly for bf.

Proof. Again consider a hypergraph with the base set Q'' and every 2-tuple of these points as a hyperedge. Color a hyperedge 'red' if the convex hull of the corresponding 2-tuple intersects 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 10, it follows that the convex hull of two of the points of these 5 points intersects 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 C. □

Lemma 14. *Given a set of points R in convex position in $\mathbb{R}^3$, $|R| \geq 5$, and a convex set C that intersects every edge spanned by the points in R, a Radon point of R is contained in C.*

Proof. Take the Radon partition of any five points in R. See Figure 5.2 (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 23. *Let P be a set of n points in $\mathbb{R}^3$. Then there exists a constant $c = f_3(g(h(5)))$ 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 $Q' \cup Q$ is a weak ϵ-net for P of size $O(\epsilon^{-5} \log^5(\epsilon^{-1}))$.*

Proof. Fix any convex set C containing at least ϵn points of P. Without loss of generality, we can assume that C is a polytope (e.g., take the convex hull of the points of P contained in C). Furthermore, one can assume that C is a full-dimensional polytope (since for a fixed weak ϵ-net Q', and each lower-dimensional polytope C' not hit by Q', there exists a full-dimensional polytope containing C' also not hit by Q').

For a large enough constant c (depending on $f_d(\cdot), g(\cdot), h(\cdot)$), by Lemma 11, Lemma 12 and Lemma 13, there exists a set of at least five points such that C intersects every edge spanned by these points. Lemma 14 then implies that Q' is a weak ϵ-net. It should be noted that the constant c has a very bad dependence on d since $f_d(\cdot)$ has a very bad dependence on d and both $g(\cdot)$ and $h(\cdot)$ are exponential functions. $\square$

Remark: In [37], 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.

5.5 Higher Dimensions

Given a finite set P of points in $\mathbb{R}^d$, the optimal weak ϵ-net can consist of any subset of $\mathbb{R}^d$. However, arguing similar to [37], 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. 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 $\mathbb{R}^d$. Consider all such points formed by the intersection of d hyperplanes (i.e. the vertex set of the hyperplanes spanned by the point set). This is the required point set, which we denote by $\Xi(P)$.

Lemma 15. *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 $\epsilon > 0$.*

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 sets [35]. Note that t lies on a facet of each of these convex sets, 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 sets containing r. The proof follows. $\square$

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

Theorem 24. *Given a set P of n points in $\mathbb{R}^d$, there is a $k(d)$ such that if Q is a random sample of size $k(d)\epsilon^{-1}\log(\epsilon^{-1})$ from P, then with constant probability, $Q' = Q \cup \Xi(Q)$ is a weak ϵ-net for P.*

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

Set $c = f_d((d+1)^2)$, where $f_d(\cdot)$ is as in Lemma 11 and set $k(d) = \lambda c$ for a large enough constant λ so that with constant probability, Q is an ϵ-net with respect to the range space $(\mathrm{P}, \mathcal{R}_c)$. Let C be any convex set containing at least ϵn points of P and assume $C \cap Q = \emptyset$. Then C cannot be separated from Q by c hyperplanes, otherwise the intersection of the halfspaces containing 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 $(\mathrm{P}, \mathcal{R}_c)$. Again assume, as in Theorem 23, that C is a full-dimensional polytope. By Lemma 11, 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 C.

By Lemma 1 of [37], Q' contains a centerpoint, say q, of the set S. We claim that q is contained in C. Otherwise, by the separation theorem, there exists a halfspace h^- containing q such that $h^- \cap C = \emptyset$. 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 C, a contradiction. □

In the above proof, we used the fact that Q' contains the centerpoint of every subset of Q. However, the proof goes through even if Q' has only a *deep-point* of every *big* subset of Q. Given a finite set S, a *deep-point* is a point $s \in \mathbb{R}^d$ such that any halfspace containing s

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 of size at least $(d+1)^2$ is present in $c(Q)$. The proof above implies the following.

Corollary 3. *If $c(Q)$ has size $O(m^t)$ for any set Q of size m, one can construct a weak ϵ-net for any point set of size $O(\epsilon^{-t}\log^t(\epsilon^{-1}))$.*

5.6 Conclusion

This chapter presented a connection between weak and strong ϵ-nets which allows the construction of weak ϵ-nets from a small number of randomly sampled input points. However, the size of the weak ϵ-net obtained this way is much larger than the best known upper bounds. It would be nice to improve the upper bound on the size of the weak ϵ-net that can be constructed from a small number of input points. The conjecture that the correct upper bound on the size of weak ϵ-nets is $O(\epsilon^{-1}\mathrm{polylog}(\epsilon^{-1}))$ remains open.

Bibliography

[1] P. K. Agarwal and N. H. Mustafa. Independent set of intersection graphs of convex objects in 2d. *Comput. Geom.*, 34(2):83–95, 2006.

[2] P. K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *J. ACM*, 51(2):139–186, 2004.

[3] P. K. Agarwal and M. Sharir. Pseudo-line arrangements: duality, algorithms, and applications. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 800–809, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[4] P. K. Agarwal, M. J. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Comput. Geom.*, 11(3-4):209–218, 1998.

[5] N. Alon, I. Bárány, Z. Füredi, and D. J. Kleitman. Point selections and weak e-nets for convex hulls. *Combinatorics, Probability & Computing*, 1:189–200, 1992.

[6] N. Alon and D. Kleitman. Piercing convex sets and the hadwiger debrunner (p; q)-problem. *Adv. Math.*, 96(1):103–112, 1992.

[7] K. Appel and W. Haken. Every planar map is four colorable. *Journal of Mathematics*, 21:108–121, 1977.

[8] B. Aronov, F. Aurenhammer, F. Hurtado, S. Langerman, D. Rappaport, S. Smorodinsky, and C. Seara. Small weak epsilon nets. In *Proc. 17th Canadian Conference on Computational Geometry*, 2005.

[9] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth. Classifying learnable geometric concepts with the vapnik-chervonenkis dimension. In *STOC '86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 273–282, New York, NY, USA, 1986. ACM.

[10] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[11] P. Carmi, M. J. Katz, and N. Lev-Tov. Covering points by unit disks of fixed location. In *ISAAC*, pages 644–655, 2007.

[12] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003.

[13] B. Chazelle, H. Edelsbrunner, M. Grigni, L. Guibas, M. Sharir, and E. Welzl. Improved bounds on weak eps-nets for convex sets. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 495–504, 1993.

[14] M. Chrobak and D. Eppstein. Planar orientations with low out-degree and compaction of adjacency matrices. *Theor. Comput. Sci.*, 86(2):243–266, 1991.

[15] K. L. Clarkson, D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng. Approximating center points with iterative radon points. *Int. J. Comput. Geometry Appl*, 6(3):357–377, 1996.

[16] K. L. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.

[17] L. Danzer. On the solution of the problem of gallai about circular discs in the euclidean plane (in german). *Stud. Sci. Math. Hung.*, 21:111–134, 1986.

[18] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational geometry: algorithms and applications*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1997.

[19] R. Diestel. *Graph theory*. Springer-Verlag, New York, 2000.

[20] D. L. Donoho and M. Gasko. Breakdown properties of location estimates based on halfspace depth and projected outlyingness. *The Annals of Statistics*, 20:1803–1827, 1994.

[21] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.

[22] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the vc-dimension is small. *Inf. Process. Lett.*, 95:358–362, 2005.

[23] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York, NY, 1979.

[24] M. T. Goodrich and E. A. Ramos. Bounded-independence derandomization of geometric partitioning with applications to parallel fixed-dimensional linear programming. *Discrete & Computational Geometry*, 18(4):397–420, 1997.

[25] J. Håstad. Clique is hard to approximate within $n^{\text{1-epsilon}}$. *Electronic Colloquium on Computational Complexity (ECCC)*, 4(38), 1997.

[26] D. Haussler and E. Welzl. ϵ-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[27] D. S. Hochbaum and W. Maass. Fast approximation algorithms for a nonconvex covering problem. *J. Algorithms*, 8(3):305–323, 1987.

[28] H. B. H. III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. Nc-approximation schemes for np- and pspace-hard problems for geometric graphs. *J. Algorithms*, 26(2):238–274, 1998.

[29] H. Imai and T. Asano. Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. *J. Algorithms*, 4(4):310–323, 1983.

[30] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm. *SIAM J. Comput.*, 15(1):287–299, 1986.

[31] J. Komlós, J. Pach, and G. J. Woeginger. Almost tight bounds for ϵ-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.

[32] J. Komlós, J. Pach, and G. J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete Computational Geometry*, pages 163–173, 1992.

[33] I. Koutis and G. Miller. A linear work, $O(n^{1/6})$ time, parallel algorithm for solving planar laplacians. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1002–1011, 2007.

[34] J. Matousek. Reporting points in halfspaces. *Comput. Geom.*, 2:169–186, 1992.

[35] J. Matousek. *Lectures in Discrete Geometry*. Springer-Verlag, New York, NY, 2002.

[36] J. Matousek, R. Seidel, and E. Welz. How to net a lot with little: Small ϵ-nets for disks and halfspaces. In *Symposium on Computational Geometry*, pages 16–22, 1990 (Corrected version: http://kam.mff.cuni.cz/ matousek/enets3.ps.gz).

[37] J. Matousek and U. Wagner. New constructions of weak ε-nets. *Discrete & Computational Geometry*, 32(2):195–206, 2004.

[38] J. Matousek, E. Welzl, and L. Wernisch. Discrepancy and approximations for bounded vc-dimension. *Combinatorica*, 13(4):455–466, 1993.

[39] J. Matoušek. Approximations and optimal geometric divide-and-conquer. *J. Comput. Syst. Sci.*, 50(2):203–208, 1995.

[40] G. Miller, S. Teng, W. Thurston, and S. Vavasis. Automatic mesh partitioning. *Workshop on Sparse Matrix Computations: Graph Theory Issues and Algorithms*, 1993.

[41] G. L. Miller, S.-H. Teng, W. P. Thurston, and S. A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *Journal of ACM*, 44:1–29, 1997.

[42] S. Narayanappa and P. Vojtechovský. An improved approximation factor for the unit disk covering problem. In *CCCG*, 2006.

[43] J. Pach and P. Agarwal. *Combinatorial Geometry*. Wiley Interscience, New York, 1995.

[44] J. Pach and G. J. Woeginger. Some new bounds for epsilon-nets. In *Symposium on Computational Geometry*, pages 10–15, 1990.

[45] E. Pyrga and S. Ray. New existence proofs for epsilon-nets. In *Proceedings of Symposium on Computational Geometry*, pages 199–207, 2008.

[46] R. Raz and M. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of STOC*, pages 475–484, 1997.

[47] N. Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13(1):145–147, 1972.

[48] S. Smorodinsky. On the chromatic number of some geometric hypergraphs. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 316–323, New York, NY, USA, 2006. ACM.

[49] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, 1971.

[50] F. F. Yao. A 3-space partition and its applications. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 258–263, 1983.