How to Net a Lot with Little: Small ε-Nets for Disks and Halfspaces*

Jiří Matoušek
Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, CSFR
and
Institut für Informatik, Freie Universität Berlin
Arnimallee 2-6, W 1000 Berlin 33, Germany
e-mail: matousek@tcs.fu-berlin.de

Raimund Seidel
Computer Science Division, University of California Berkeley
Berkeley CA 94720, USA
e-mail: seidel@cs.berkeley.edu

Emo Welzl
Institut für Informatik, Freie Universität Berlin
Arnimallee 2-6, W 1000 Berlin 33, Germany
e-mail: emo@tcs.fu-berlin.de

July 6, 2000

Abstract

It is known that in general range spaces of VC-dimension $d > 1$ require ε-nets to be of size at least $\Omega(\frac{d}{\varepsilon}\log\frac{1}{\varepsilon})$. We address the question whether this general lower bound is valid for the special range spaces that typically arise in computational geometry. We show that disks and pseudo-disks in the plane as well as halfspaces in R^3 allow ε-nets of size only $O(\frac{1}{\varepsilon})$, which is best possible up to a multiplicative constant. The analogous questions for higher-dimensional spaces remain open.

1 Introduction

A *range space* is a pair $(X, \mathcal{R})$, where X is a set and $\mathcal{R}$ is a family of subsets of X. We usually refer to the members of X as *points* and the members of $\mathcal{R}$ as *ranges*. The *Vapnik-Chervonenkis dimension*, or, *VC-dimension*, of a range space $(X, \mathcal{R})$ is defined to be the cardinality of the largest subset A of X for which $\{r \cap A | r \in \mathcal{R}\}$ is the powerset of A. As an example consider $(\mathrm{R}^d, \mathcal{H}_d)$, where $\mathcal{H}_d$ is the family of all closed halfspaces in R^d. Using Radon's theorem it can easily be shown that $(\mathrm{R}^d, \mathcal{H}_d)$ has VC-dimension $d+1$.

*Large Portions of the research reported here were conducted while R.S. visited the Free University of Berlin, and while all three authors visited DIMACS at Princeton University. J.M. and E.W. acknowledge support by the Deutsche Forschungsgemeinschaft under Grant We 1265/1–2, Schwerpunktprogramm "Datenstrukturen und effiziente Algorithmen", and R.S. acknowledges support by NSF grant CCR-8809040.

The notions of range space and VC-dimension were inspired by the work of Vapnik and Chervonenkis [VC71], and they were introduced to computational geometry by Haussler and Welzl [HW87], who used it to develop linear sized data structures for certain range query problems with sublinear query time. Later Chazelle and Welzl [CW89] showed that in the so-called arithmetic model range spaces that admit such data structures are precisely those that have finite VC-dimension.

An important notion related to range spaces of finite VC-dimension is an *ε-net*. Let $T \subseteq X$ be a finite set, and ε a real parameter, $0 < \varepsilon \leq 1$. A set $N \subseteq T$ is called an *ε-net for T with respect to $\mathcal{R}$* iff for every range $r \in \mathcal{R}$ with $|r \cap T| > \varepsilon|T|$ we have $r \cap N \neq \emptyset$. Haussler and Welzl proved that in any range space $(X, \mathcal{R})$ any finite $T \subset X$ has an ε-net N of size $O(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon})$ for every ε. As a matter of fact, they showed that a random subset of T of this size is an ε-net for T with probability close to 1. Note that the size of N depends on d and ε only and not on the size of T. The upper bound for the size of ε-nets was later improved by Blumer et al. [BEHW89] to $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$, and to $(1 + o(1))\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}$ by Komlós, Pach, and Wöginger [KPW91].

The relative crudeness of the proof method and the observation that a number of simple range spaces admit ε-nets of size $O(\frac{d}{\varepsilon})$ suggested the possibility that the $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$ upper bound might admit improvement and that in fact range spaces of VC-dimension d always allow ε-nets of size $O(\frac{d}{\varepsilon})$. This hope was shattered by Komlós et al. [KPW91], who showed that for every $d > 1$ there exist range spaces of VC-dimension d for which ε-nets must have size at least $\Omega(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$. Ironically, also this existence proof is based on randomization. Choosing parameters n, r, and p judiciously, Komlós et al. show that with high probability the range space $(X, \mathcal{R})$, where $|X| = n$ and $\mathcal{R}$ just contains r-subsets of X, each included with independent probability p, has VC-dimension at most d but requires ε-nets of size $\Omega(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$.

However, this still leaves open the possibility that the range spaces that typically arise in a geometric context have, in comparison to the randomness of the spaces of [KPW91], sufficient inherent structure so that ε-nets of size $O(\frac{d}{\varepsilon})$ are possible for them after all. This paper addresses that question and provides some partial answers.

One of the first applications of ε-nets and related probabilistic techniques (see e.g., [Cla87]) was in the following problem: Given a collection H of n hyperplanes in R^d and a parameter r, $1 < r \leq n$, subdivide the space into a possibly small number of simplices in such a way that the interior of every simplex is intersected by at most n/r hyperlanes of H. Later various authors managed to remove the extra logarithmic factors introduced by a direct use of ε-nets ([CF90], [Mat90], [Aga90], [Mat91a], [Cha91]). However, these results seem to have nothing to do with improving the bounds for the size of ε-nets in geometric settings.

Our main interest are the range spaces $(\mathrm{R}^d, \mathcal{H}_d)$ with $d > 2$. This special attention is warranted since due to an observation of Frances and Andy Yao [YY85] many finite-dimensional geometric range spaces can be embedded in $(\mathrm{R}^d, \mathcal{H}_d)$ for sufficiently large d. Take, for example, the range space $(\mathbf{H}_2, \mathcal{L}_2)$, where $\mathbf{H}_2$ denotes the set of all nonvertical lines in the plane, and every line segment in the plane defines the range of all lines intersecting it. Observe now, that for a line segment with endpoints (a_1, a_2) and (b_1, b_2), the lines $y = kx + d$ which intersect this line segment are exactly those which satisfy the polynomial inequality $p(k,d) := (ka_1 + d - a_2)(kb_1 + d - b_2) \leq 0$. So if we map lines $y = kx + d$ in the plane to points (k^2, kd, d^2, k, d) in R^5, then, for every line segment, there is a halfspace in R^5 that contains exactly all points that come from lines that intersect this line segment. In a dual setting, this corresponds to an embedding of the range space defined by double wedges in the plane into a range space $(\mathrm{R}^d, \mathcal{H}_d)$ (with $d = 5$). On the other hand, we could not find any embedding of the range space defined by *wedges* (i.e. intersections of two halfplanes) in the plane into some $(\mathrm{R}^d, \mathcal{H}_d)$, and we conjecture that none exists.

It has been known for some time that $(\mathrm{R}^2, \mathcal{H}_2)$ admits ε-nets of size $O(\frac{1}{\varepsilon})$. Pach and Wöginger [PW] even gave an $\lceil \frac{2}{\varepsilon} \rceil - 1$ upper bound for such ε-nets and showed this exact bound to be tight within an additive constant of one. We show in this paper that $(\mathrm{R}^3, \mathcal{H}_3)$ also admits ε-nets of size $O(\frac{1}{\varepsilon})$. Unfortunately, the question remains open for $d > 3$. However, it is interesting to observe that points on the moment curve, which often form an extreme configuration for geometric combinatorial problems, do allow ε-nets of size $O(\frac{d}{\varepsilon})$ with respect to halfspaces.

Our proof is constructive. It reduces the problem of finding an ε-net for a set of n points in R^3 with respect to halfspaces to the problem of finding such an ε-net for n points in R^3 in convex position. This problem in turn is reduced to the problem of finding an ε-net for an n-point set in the plane with respect to pseudo-disks, which we then solve directly. Although algorithms are not our major concern, our proof yields a deterministic polynomial time algorithm of worst case complexity at most $O(n^3)$.

Section 2 describes the reductions. Auxiliary results about pseudo-disks are proven in Section 3. Section 4 outlines the construction of ε-nets for planar point sets with respect to pseudo-disks. Finally, Section 5 briefly discusses the question of so-called weak ε-nets, where the points in the net can be arbitrary and need not be part of the input set. We show that this question admits a straightforward answer for all dimensions.

Let us briefly summarize recent new developments. After a preliminary version of the present paper appeared, Clarkson outlined an alternative probabilistic proof of our main result (concerning ε-nets with respect to halfspaces in dimension 3) in an unpublished note [Cla89]. A similar proof was given in [Mat91b], where the result easily follows from a more general statement. Concerning weak ε-nets, interesting new results were attained by Alon et al. [ABFK91], and they were applied by Alon and Kleitman [AK91] for a solution of an old problem from combinatorial geometry.

2 The Reductions

A finite set $T \subset \mathrm{R}^3$ is said to be in *convex position* iff every point of T is a vertex of the convex hull of T. The following lemma shows that as far as ε-nets with respect to halfspaces are concerned it suffices to consider only point sets in convex position (this is true also in higher dimensions R^d, and the proof of the following lemma also works there with 'd' instead of '3').

Lemma 2.1 *Let $0 < \varepsilon \leq 1$. If for every finite set $T \subset \mathrm{R}^3$ in convex position there is an ε-net with respect to halfspaces of size $s(\varepsilon)$, then for every finite set in R^3 there exists an ε-net with respect to halfspaces of size at most $3s(\varepsilon)+1$.*

Proof: Let $S \subset \mathrm{R}^3$ contain n points. We want to construct an ε-net for S of size at most $3s(\varepsilon)+1$, i.e. we want to find a subset $N \subset S$ of this size so that whenever a halfspace contains more than εn points of S it also contains a point of N.

We may assume that S has no four points co-planar, as it is always possible to perturb S into such a non-degenerate position so that the collection of subsets of S of the form $S \cap H$, where H is a halfspace, does not decrease [EW85]. An ε-net for the perturbed set then yields an ε-net for the original set.

Let V be the set of vertices of the convex hull of S and I be the remaining points of S. Choose a vertex $v \in V$ and for each point $p \in I$ let ℓ_p be the line through v and p. By non-degeneracy all the lines are distinct and each line ℓ_p meets the boundary of conv$S$ besides in v in the relative interior of some facet F_p. Now, for each $p \in I$ choose a point p' on ℓ_p "just outside" of F_p, i.e. $p' \notin \mathrm{conv}S$ and F_p is the only facet visible from p'. Moreover, do this in such a way so that $S' = V \cup \{p' | p \in I\}$ is in convex position. This is clearly always possible.

Let N' be an ε-net for S'. By assumption it has size at most $s(\varepsilon)$. Now form the set $N \subset S$ as follows: include v; for each $x \in N'$, if $x \in V$ then include x, otherwise $x = p'$ for some $p \in I$ and in this case include the three vertices of the facet F_p (which must be a triangle by non-degeneracy). Clearly N has size at most $3s(\varepsilon)+1$. We claim that it is an ε-net for S.

Let H be a halfspace with $|H \cap S| > \varepsilon n$. We need to show that it contains a point of N. This is trivially the case if $v \in H$. So assume $v \notin H$. In that case clearly $|H \cap S'| \geq |H \cap S| > \varepsilon n$, and therefore H contains some point $x \in N'$. If $x \in V$ then $x \in N$ and we are done. Otherwise $x = p'$ for some $p \in I$ and in this case H must contain at least one of the vertices of F_p. But they are all in N. ◻

Let us call a compact convex set in the plane with smooth boundary of positive curvature an *oval*. We call a family $\mathcal{D}$ of ovals a *family of pseudo-disks* iff for every three non-collinear points x, y, z in the plane there is a unique oval $D = D_{xyz} \in \mathcal{D}$ that has x, y and z on its boundary. Examples of families of pseudo-disks are ordinary disks, or the set of all homothets of a fixed oval. We show in Section 3 that combinatorially pseudo-disks behave in many ways like ordinary disks, in particular the boundary of any two intersect in at most two points. It is also not hard to prove that for any family of pseudo-disks $\mathcal{D}$, the range space $(\mathrm{R}^2, \mathcal{D})$ has VC-dimension 3.

We say that a finite set S in the plane is *non-$\mathcal{D}$-degenerate* iff no three points of S are collinear and no four points of S lie on the boundary of the same oval $D \in \mathcal{D}$.

Lemma 2.2 *Let* $0 < \varepsilon \leq 1$. If for every family $\mathcal{D}$ of pseudodisks and for every non-$\mathcal{D}$-degenerate finite point set in the plane a there exists an ε-net of size at most $t(\varepsilon)$ with respect to $\mathcal{D}$, *then for every finite set in* R^3 *in convex position there exists an* ε*-net of size at most* $t(\varepsilon) + 1$ *with respect to halfspaces.*

Proof: Let T be a set of n points in R^3 in convex position and let K be a compact convex body in R^3 whose boundary B is smooth, has positive curvature everywhere, and contains T. Clearly such a K always exists. Let c be a point in T, let h_c be the tangent plane of K at c, and let h be the unique other tangent plane of K that is parallel to h_c. There is a natural bijection ϕ from $B \setminus \{c\}$ to h, mapping a point $x \in B \setminus \{c\}$ to the intersection point of the plane h with the straight line through c and x (ϕ is thus a "pseudo-stereographic projection"). If H is a halfspace in R^3 intersecting the interior of K, but $c \notin H$, then ϕ maps $H \cap B$ to an oval in h. Moreover, the collection $\mathcal{D}$ of such ovals forms a family of pseudo-disks in h. (The unique oval with non-collinear $x, y, z \in h$ on its boundary is given by $\phi(H \cap B)$, where H is the unique halfspace that has $\phi^{-1}(x), \phi^{-1}(y), \phi^{-1}(z)$ on its boundary and does not contain c.) Now let $T' = \phi(T \setminus \{c\})$. As we remarked above, we may assume that T is in general position, and then we get that T' is non-$\mathcal{D}$-degenerate. By assumption there is an ε-net $N' \subset T'$ of size at most $t(\varepsilon)$ for T' with respect to $\mathcal{D}$. By construction, clearly $N = \phi^{-1}(N') \cup \{c\}$ is an ε-net for T with respect to halfspaces, and $|N| \leq t(\varepsilon) + 1$. $\square$

Lemma 2.3 *Let* $0 < \varepsilon \leq 1$. *If for every finite set in the plane there exists an* ε*-net of size at most* $t(\varepsilon)$ *with respect to a family of pseudo-disks* $\mathcal{D}$, *then for every finite set in* R^3 *there exists an* ε*-net of size at most* $3t(\varepsilon) + 1$ *with respect to halfspaces.*

Proof: Combining Lemmas 2.1 and 2.2 immediately yields the result with a bound of $3t(\varepsilon) + 4$. The slight improvement to $3t(\varepsilon) + 1$ can be obtained by identifying the vertices v and c of the constructions in the proofs. $\square$

3 Properties of Pseudo-Disks

In this section we establish some properties of families of pseudo-disks needed in the subsequent ε-net construction. For ordinary circular disks, these properties are mostly trivial or well-known.

From now on S will denote a non-$\mathcal{D}$-degenerate set, $\mathcal{D}$ will stand for a fixed family of pseudo-disks, and the word "disk" will mean a pseudo-disk from $\mathcal{D}$.

Claim 3.1 *The boundaries of two distinct disks intersect in at most two points.*

Proof: If they intersected in three points these three points would not define a unique disk. $\square$

Claim 3.2 *Let* D_1, D_2 *be disks with boundaries* C_1,C_2, *respectively. If* C_1 *and* C_2 *intersect at two points* p *and* q, then they cross at p and q.

Proof: Suppose not, and w.l.o.g. assume $D_1 \subset D_2$. Choose a point $x_1 \notin \{p, q\}$ on C_1 and a point $x_2 \notin \{p, q\}$ on C_2 such that x_1 and x_2 do not lie on the boundary of the same connected component of $D_2 \setminus C_1$. With this choice the points p, x_1, and x_2 cannot be collinear and thus they define a unique disk D. But now C, the boundary of D must intersect either C_1 or C_2 in more than two points in contradiction to Claim 3.1. Why? The curve C yields two distinct paths from x_1 to x_2. One of the paths contains p, which counts as two intersections (both for C_1 and C_2). As x_1 and x_2 do not lie on the boundary of the same connected component of $\mathrm{R}^2 \setminus (C_1 \cup C_2)$ the other path also must contain a point of $C_1 \cup C_2$. But together with points x_1 and x_2 this gives a total of five intersections in $C \cap C_1$ and $C \cap C_2$. □

We call two points $p, q \in S$ *$\mathcal{D}$-Delaunay neighbors*[1] iff there exists a disk D that has p and q on its boundary but no point of S in its interior. The *$\mathcal{D}$-Delaunay graph of S*, for short $\mathcal{D}$-$DT(S)$ is a graph embedded in the plane with vertex set S and straight edges connecting $\mathcal{D}$-Delaunay neighbors. We call these edges $\mathcal{D}$-Delaunay edges.

Claim 3.3 *Let e be a $\mathcal{D}$-Delaunay edge joining $\mathcal{D}$-Delaunay neighbors $p,q \in S$ and let D be a disk such that e intersects* $\mathrm{int}D$*, but $p, q \notin \mathrm{int}D$. At least one of the two components of $D \setminus e$ contains no point of S.*

Proof: Let D_e be a disk with p and q on its boundary but no point of S in its interior. By convexity $e \subset D_e$, and hence D and D_e have non-empty intersection.

If $D = D_e$, then the claim follows as (by non-$\mathcal{D}$-degeneracy) at most three points of S lie on the boundary of D.

If $D \subset D_e$, then at least one component of $D \setminus e$ is contained in $\mathrm{int}D_e$ and hence contains no point of S.

The case D_e a proper subset of D is impossible as a consequence of Claim 3.2.

In the remaining possible case Claim 3.2 implies that $D \setminus \mathrm{bd}D_e$ consists of two components, one of which, call it F, is convex and lies in D_e. The edge e intersects F, and since $p, q \notin \mathrm{int}D$, it follows that $F \setminus e$ consists of two components, of which one is also a component of $D \setminus e$. But that component is contained in $\mathrm{int}D_e$ and hence contains no point of S. □

Claim 3.4 *Two $\mathcal{D}$-Delaunay edges e and f can only intersect in a common endpoint. In other words, $\mathcal{D}$-$DT(S)$ is a planar graph.*

Proof: Clearly an endpoint of one edge cannot be contained in the relative interior of the other. Thus if there is an intersection point x that is not an endpoint of both edges, it must lie in the relative interior of each edge. For this case let D be a disk that has the endpoints of e on its boundary but no point of S in it interior. The endpoints of f cannot be contained in D. But f intersects $\mathrm{int}D$ since $x \in \mathrm{relint}\, e \subset \mathrm{int}D$. But now $D \setminus f$ has a point of S in each component (an endpoint of e), in contradiction to Claim 3.3. □

[1] non-$\mathcal{D}$-generate? $\mathcal{D}$-launay?

Claim 3.5 *If S is a finite set of $n > 2$ points in R^2 in non-$\mathcal{D}$-degenerate position, then $\mathcal{D}$-$DT(S)$ is a triangulation of S.*

Proof: We need to show (i) that every edge of the convex hull of S is part of a triangle in $\mathcal{D}$-$DT(S)$, (ii) that every $\mathcal{D}$-Delaunay edge internal to the convex hull is part of two triangles of $\mathcal{D}$-$DT(S)$, and (iii) that every point of S is a vertex of $\mathcal{D}$-$DT(S)$.

First consider (ii): Let e be the $\mathcal{D}$-Delaunay edge with endpoints p and q, which does not lie on the boundary of the convex hull of S. Thus the set $T = S \cap H$, where H is an open halfplane with e on its boundary is non-empty. By non-$\mathcal{D}$-degeneracy of S all the disks D_{pqx} with $x \in T$ must be distinct. As the boundaries of all these disks intersect in p and in q Claim 3.2 implies that $\bigcap\{D_{pqx} \cap H | x \in T\}$ equals $D_{pqr} \cap H$ for some $r \in T$. Note that $D_{pqr} \cap H$ contains no point of T and hence of S in its interior. However, $D_{pqr} \cap H$ does contain r (on its boundary). Claim 3.3 now implies that D_{pqr} cannot contain any point of S in its interior. Thus p, r and q, r are $\mathcal{D}$-Delaunay neighbors and therefore p, q, r is a triangle in $\mathcal{D}$-$DT(S)$. Applying this reasoning to both possible choices of H yields two triangles for the edge e.

The same reasoning also applies for proving (i). However, only one choice is possible for H so that $H \cap S \neq \emptyset$. Moreover, Claim 3.3 need not be used, as the complement of H only contains the points p and q. Note that the fact that the obtained disk D_{pqr} has no points in its interior establishes that the convex hull edge e is a $\mathcal{D}$-Delaunay edge, which was not clear a priori.

(i) and (ii) imply that the union of the triangles of $\mathcal{D}$-$DT(S)$ covers the convex hull of S. Thus any point $x \in S$ has to be contained in some triangle p, q, r of $\mathcal{D}$-$DT(S)$ and hence x is contained in D_{pqr}. It cannot be contained in the interior of D_{pqr} and thus it must be one of the corners of the triangle p, q, r, which proves (iii). $\square$

Claim 3.6 *Let D be a disk. The subgraph G of $\mathcal{D}$-$DT(S)$ that is induced by the points of S that lie in D is connected.*

Proof: Assign to each connected component of G a unique color and "paint" each triangle of $\mathcal{D}$-$DT(S)$ that is incident to one of the vertices of G with the color of that vertex. By the convexity of D, no triangle can be painted twice. The union of the triangles painted the same color is connected and its boundary consists of $\mathcal{D}$-Delaunay edges that either lie outside of D or cut across D. If G has more than one component there must be such a $\mathcal{D}$-Delaunay edge that cuts across D, but vertices of G, i.e. points of S in D, lie on both of its sides. But this contradicts Claim 3.3 and thus G cannot have more than one component. $\square$

Claim 3.7 *Let S be a finite point set.*

(i) *For any disk D, there exists a disk D' with $S \cap D' = S \cap \mathrm{int} D' = S \cap D$.*

(ii) *For any disk D, there exists a disk D' with $S \cap D' = S \cap \mathrm{int} D' = S \cap \mathrm{int} D$.*

Proof: (i) Consider a closed halfplane h^+ with $h^+ \cap \mathrm{int}D \neq \emptyset$ and $h^+ \cap D \cap S = \emptyset$. The line h bounding h^+ intersects the boundary of D in two points x and y. Every point p in $S \cap h^-$ (with h^- the complement of h^+) defines a unique disk D_{xyp} with points $\{x, y, p\}$ on its boundary. The sets $D_{xyp} \cap h^-$ are ordered by inclusion. For each point p in $(S \cap h^-) \setminus D$, the disk D_{xyp} contains $D \cap S$ and $h^+ \cap D_{xyp} \cap S = \emptyset$. Let now p_0 be a point in $(S \cap h^-) \setminus D$ with $D_{xyp_0} \cap h^-$ minimal. For an arbitrary point $q \in D_{xyp_0} \setminus D$, the disk D_{xyq} satisfies $S \cap D_{xyq} = S \cap \mathrm{int}D_{xyq} = S \cap D$.

The proof for (ii) proceeds similarly. □

Claim 3.8 *Let S be a finite set in non-$\mathcal{D}$-degenerate position, let (p, q) be a $\mathcal{D}$-Delaunay edge in $\mathcal{D}$-$DT(S)$. Then there exists a disk D with p, q on the boundary and with $D \cap S = \{p, q\}$.*

Proof: By definition of a $\mathcal{D}$-Delaunay edge, and because of non-$\mathcal{D}$-degenerate position, there is a disk D with $\mathrm{int}D \cap S = \emptyset$, and $D \cap S = \{p, q\}$ or $D \cap S = \{p, q, r\}$ for some third point r. In the first case nothing is to prove. In the second case, we can proceed as in the proof of the previous claim, but choosing p and q for x and y. □

4 Constructing ε-Nets for Pseudo-Disks

In this section we prove the following:

Theorem 1 *Let $0 < \varepsilon \leq 1$ and let $\mathcal{D}$ be a family of pseudo-disks. For every finite point set S* in non-$\mathcal{D}$-degenerate position *in the plane there exists an ε-net with respect to $\mathcal{D}$ of size $O(1/\varepsilon)$.*

Together with Lemma 2.3 this yields the main result of this paper:

Theorem 2 *Let $0 < \varepsilon \leq 1$. For every finite set T in the R^3 there exists an ε-net with respect to halfspaces of size $O(1/\varepsilon)$.*

Let us begin with the proof of Theorem 1. For a better intuition, the reader may imagine that we talk about ordinary circular disks.

Let $0 < \varepsilon \leq 1$ and let S be a set of $n > 2$ points in the plane. We may assume S to be in non-$\mathcal{D}$-degenerate position. We are to determine a subset N of S of size $O(1/\varepsilon)$ so that any disk of $\mathcal{D}$ that contains εn points of S also contains a point of N.

Let $\delta = \varepsilon/6$. Let $S_1, \ldots, S_j$ be disjoint subsets of S satisfying the following properties: their union contains all vertices of $\mathrm{conv}S$; each S_i is representable as $S \cap H_i$ for some halfplane H_i; the cardinality of S_i is $\lfloor \delta n \rfloor$ for $1 \leq i < j$, and $|S_j| \leq \lfloor \delta n \rfloor$. Such a collection of subsets of S can easily be constructed by repeatedly "biting off" subsets of S with halfplanes. Note that each S_i can also be represented as $S \cap D_i$ for some (suitably large) disk D_i. Next let $S_{j+1}, \ldots, S_k$ be a maximal collection of disjoint subsets of $S \setminus \bigcup_{1 \leq i \leq j} S_i$ satisfying both $S_i = S \cap D_i$ for some disk D_i and $|S_i| = \lfloor \delta n \rfloor$ for $j < i \leq k$. Since all S_i's are disjoint and all but maybe one contain $\lfloor \delta n \rfloor$ points of S, clearly k is at most $\lceil 1/\delta \rceil + 1$.

For each i, $1 \leq i \leq k$, let us "color" the points of S_i with color i, and let us call D_i the *disk defining color* i. Let $\bar{S}$ be the set of all colored points and let us call the points in $S \setminus \bar{S}$ *colorless.*

Claim 4.1 *Any disk $D \in \mathcal{D}$ containing no colored points contains less than $\lfloor \delta n \rfloor$ points of S. Also, if* $\mathrm{int}D$ *contains no colored points, then* $|\mathrm{int}D \cap S| < \lfloor \delta n \rfloor$.

Proof: The statement for closed disks follows from the construction of the colored points, and the claim for open disks follows from Claim 3.7. □

Let DT be the $\mathcal{D}$-Delaunay triangulation of the set of colored points $\bar{S}$. Since each vertex of conv$S$ is colored, DT is a triangulation of conv$S$ by Claim 3.5.

We distinguish the triangles in this triangulation according to the number of colors appearing in their vertex sets, into uni-colored, bi-colored and tri-colored triangles. Similarly we speak about uni-colored and bi-colored edges. We call a maximal connected chain of bi-colored triangles in DT sharing bi-colored edges a *corridor*. We consider a bi-colored edge that is sandwiched between two tri-colored triangles to be a degenerate corridor.

Claim 4.2 *There are at most $3k - 6$ corridors in DT.*

Proof: Contract all uni-colored edges of DT, and for every collection of bi-colored multiple edges originating from the same corridor, remove all but one. The resulting planar graph G has k vertices and it may still have multiple edges, but it has only 3-sided interior faces. Such a graph has at most $3k - 6$ edges, and each edge of G corresponds to a corridor in DT. □

Since any uni-colored triangle is contained in the corresponding color-defining disk D_i, all colorless points are contained in the corridors. Break each corridor C into a minimum number of "subcorridors," i.e. subchains of the chain of triangles that form C, so that each subcorridor contains at most $\lfloor \delta n \rfloor$ colorless points. Since there are fewer than n colorless points and because of Claim 4.2 the total number of subcorridors will be $O(1/\delta)$.

Some terminology for subcorridors: Since a subcorridor is a connected chain of bi-colored triangles sharing bi-colored edges it is bounded by two uni-colored chains of edges, which we call the *sides*, and two bi-colored edges, which we call the *ends* of the subcorridor. We call the endpoints of the sides the *corners* of the subcorridor. Note that one of the sides can degenerate into a single point, in which case the subcorridor has 3 corners; otherwise it has four. (Of course, a subcorridor can also have only two corners: if it is a degenerate corridor formed by a single bi-colored edge.)

Let N_0 be the set of the corners of all subcorridors. Clearly $N_0 \subset S$. Moreover, since each subcorridor has at most 4 corners and since there are $O(1/\delta)$ subcorridors, the size of N_0 is $O(1/\delta)$, which is $O(1/\varepsilon)$, as δ is defined to be $\varepsilon/6$. We could prove that N_0 itself is an ε-net for S with respect to disks, but in order to simplify the argument, we choose an $(\varepsilon/2)$-net N_1 for S with respect to halfplanes, with $N_1 \subseteq \bar{S}$, $|N_1| = O(1/\varepsilon)$, and we set $N = N_0 \cup N_1$. We are going to show that N is an ε-net with respect to disks of $\mathcal{D}$.

Let us begin by a topological lemma.

Lemma 4.3 *Let D be an open, simply connected subset of the plane (see Fig 1). Let C be an open region bounded by a closed Jordan curve Γ, consisting of finitely many straight segments. Let $y_1, \ldots, y_k$ be points on Γ, numbered in the order of their appearance along Γ, let γ_i denote the arc of Γ between the points y_i and y_{i+1} ($y_{k+1} = y_1$). Let $B_1, \ldots, B_k$ be closed sets with $\gamma_i \subseteq B_i$ and such that $D \setminus B_i$ is connected. Finally let $x_1, \ldots, x_k$ be (not necessarily distinct) points all belonging to one connected component of $C \cap D$, and such that $x_i \notin B_i$. Then D is contained in the union $C \cup B_1 \cup \ldots \cup B_k$.*

Proof: Assume that the conclusion if false, so there exists a point t outside the union $C \cup B_1 \cup \ldots \cup B_k$. Let $\pi_1 \subset D$ be a path connecting x_1 to t; we may assume that it consists of a finite number of straight segments. Since x_1 lies inside Γ and t outside Γ, the path π_1 must intersect Γ an odd number of times[2], and hence there is some arc γ_i intersected by π_1 an odd number of times. Consider now a path $\xi \subset C \cap D$ connecting x_i to x_1, and let π_i be the path from x_i to t arising by appending π_1 to ξ. Any path $\pi \subset D$ connecting x_i to t is homotopic to π_i (inside D), and since a continuous deformation can only change the number of intersections of a path inside D with γ_i by an even number[3], such a path must intersect γ_i, and thus also B_i. This means that x_i and t lie in different components of $D \setminus B_i$, contradicting the connectedness assumption. $\square$

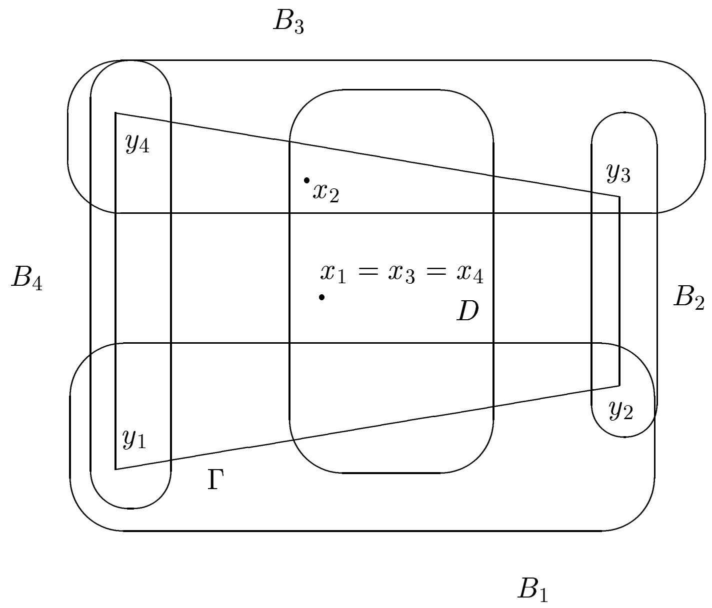

Figure 1: The topological lemma

[2]Since both Γ and π_1 consist of a finite number of segments, it is straightforward to define the "number of intersections" consistently (even if π_1 and Γ have segments in common). This is the reason we work with such paths and curves: to avoid unnecessary technical complications.

[3]We may assume that also all intermediate paths in the homotopy consist of a finite number of segments.

We are going to finish the proof of Theorem 1. By Claim 3.7, it suffices to consider an open disk D (the interior of some disk from $\mathcal{D}$) containing no point of N. If D contains no colored point, the we are done by Claim 4.1. Let us first assume that D contains colored points of at least two colors. Then, by Claim 3.6, D contains a bi-colored Delaunay edge $\{z_1, z_2\}$; let us assume that z_1 has color 1 and z_2 color 2. Let D_1, D_2 be the respective color-defining disks.

The edge $\{z_1, z_2\}$ belongs to a unique subcorridor C. Since none of z_1, z_2 are corners of C, C has four corners; let us denote them a_1, a_2, b_2, b_1 in such a way that $\{a_1, a_2\}$ and $\{b_1, b_2\}$ are the ends of C, $a_1, b_1 \in D_1$ and $a_2, b_2 \in D_2$ (Fig. 2). Further, let D_a (resp. D_b) denote a (closed) $\mathcal{D}$-Delaunay disk belonging to the edge $\{a_1, a_2\}$ (resp. $\{b_1, b_2\}$), such that D_a contains no colored points but a_1, a_2 and similarly for D_b (such disks are guaranteed by Claim 3.8). In particular, $z_1 \notin D_a, D_b$.

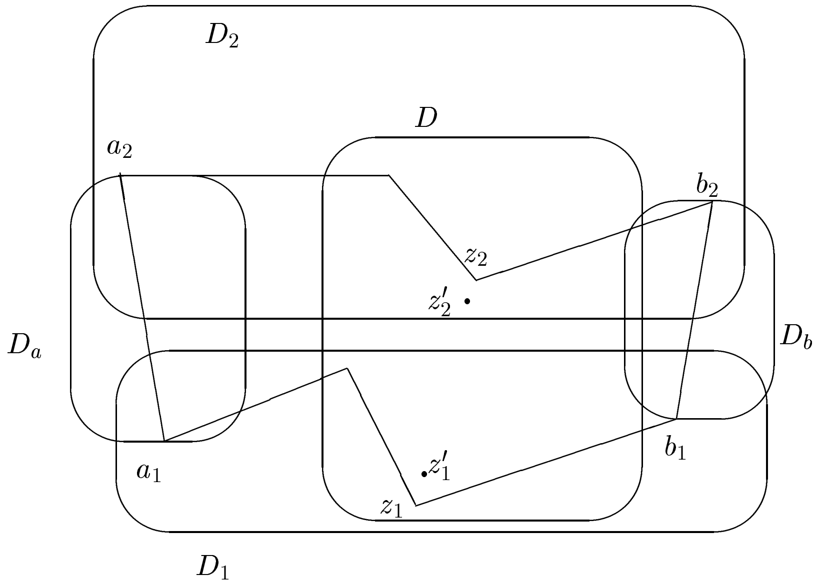

Figure 2: The subcorridor C

We observe that the open segment z_1z_2 is contained in the interior of the subcorridor C. We choose a point z_1' on that segment very close to z_1, in such a way that $z_1' \in D$, $z_1' \notin D_2, D_a, D_b$ (such a choice is possible, since z_1 has these properties, D is open and D_2, D_a, D_b are closed). Similarly, we choose a point z_2' on the segment z_1z_2 with $z_2' \in D$, $z_2' \notin D_1$. We are going to apply Lemma 4.3 to show that D is contained in the union of C, D_1, D_2, D_a and D_b. The open disk D will play the role of D in the Lemma, C will be the interior of our subcorridor, Γ will be its boundary. We let $k = 4$, $y_1 = a_1$, $y_2 = a_2$, $y_3 = b_2$, $y_4 = a_2$, $B_1 = D_a$, $B_2 = D_2$, $B_3 = D_b$, $B_4 = D_1$, $x_1 = x_2 = x_3 = z_1'$, $x_4 = z_2'$. One easily verifies that the conditions of Lemma 4.3 are satisfied, and hence we may conclude that $D \subseteq C \cup D_1 \cup D_2 \cup D_a \cup D_b$. Now $|D_1 \cap S| = |D_2 \cap S| = \lfloor \delta n \rfloor$,

$|C \cap S| \leq \delta n$ by the construction of the subcorridors. The interiors of D_a and D_b contain at most δn points of S by Claim 4.1, and therefore $|D \cap S| \leq 5\delta n < \varepsilon n$.

It remains to consider the case when D contains colored points of only one color, say color 1. If D is contained in the color-defining disk D_1, then we are done, so suppose it is not the case, and let p, q be the intersections of the boundaries of D and of D_1. By a slight perturbation of D (which does not change $S \cap D$), we can achieve that $S \cup \{p, q\}$ is in a non-$\mathcal{D}$-degenerate position. Since the boundary of a disk has at most two intersections with a line, one of the halfplanes, h, determined by the line pq, contains $D \setminus D_1$.

Let T be the set of all colored points in $h \setminus D$. If $T = \emptyset$, then $h \cap N = \emptyset$, and since N is an $\varepsilon/2$-net for halfplanes, $|D \cap S| \leq |h \cap S| + |D_1 \cap S| \leq \varepsilon n/2 + \delta n \leq \varepsilon n$. Let us consider the case $T \neq \emptyset$.

It is easily seen that for $x, y \in T$, the sets $D_{pqx} \cap h$ and $D_{pqy} \cap h$ are comparable by inclusion, and that this defines a linear order on T (the antisymmetry is obtained from the assumed non-$\mathcal{D}$-degenerate position of $S \cup \{p, q\}$). Let $z \in T$ be the minimal element in this ordering, and let $D' = \mathrm{int} D_{pqz}$. Then we have $D \subseteq D' \cup D_1$, and it suffices to show that D' contains at most $\varepsilon n - \delta n$ points of S (see Fig. 3).

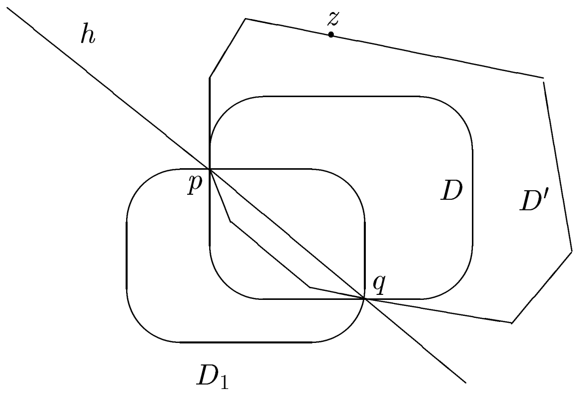

Figure 3: The proof for the uni-colored case

We note that D' may contain only colorless points and points of color 1. If it contains no colored points at all, then we are done by Claim 4.1. Otherwise let z_1 be a point of color 1 in D' such that $\{z, z_1\}$ is a $\mathcal{D}$-Delaunay edge (there must be such an edge by the connectedness of $DT \cap D_{pqz}$, see Claim 3.6). Since z has a different color than 1 (2, say), and since $z_1 \notin N$, this edge belongs to a unique subcorridor C. Further we proceed very similarly to the above bi-colored case. Let a_1, a_2, b_1, b_2 be the corners of C ($a_1 a_2$ and $b_1 b_2$ are the ends, $a_j, b_j \in D_j$, $j = 1, 2$); this time it might happen that $z = a_2$, $z = b_2$ or even $z = a_2 = b_2$. We choose a point z_1' on the segment zz_1 with $z_1' \in C$, $z_1' \in D'$, $z_1' \notin D_2, D_a, D_b$ (D_a, D_b are $\mathcal{D}$-Delaunay disks for the edges $\{a_1, a_2\}$ and $\{b_1, b_2\}$, as in Claim 3.8). We choose a point z' on the segment zz' with $z' \in D'$, $z' \in C$, $z' \notin D_1$ (we use the fact that the relative interior of the segment zz_1 is contained in D'). We now apply Lemma 4.3, and we conclude that $D' \subseteq C \cup D_1 \cup D_2 \cup D_a \cup D_b$. This means $|D' \cap S| \leq 5\delta n$, and $|D \cap S| \leq 6\delta n = \varepsilon n$. This concludes the proof of Theorem 1. $\square$

Note that in principle the ε-net construction presented in this paper can be transformed into a deterministic algorithm that runs in polynomial time, $O(n^3)$ at worst. However, we certainly would not advocate this algorithm as being practical. We find the resulting constant of proportionality also not particularly flattering.

Our construction relies heavily on planarity arguments. We see little hope in extending this type of approach to the higher-dimensional case, where it remains open whether ε-nets of size $o(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$ exist with respect to halfspaces.

5 Weak ε-Nets

The reader might wonder about the definition of ε-nets, whether it would not be simpler to work with the following notion of a *weak ε-net* : For a finite set S of points in a range space $(X, \mathcal{R})$ a set $N \subset X$ is a weak ε-net for S with respect to $\mathcal{R}$ iff all $r \in \mathcal{R}$ with $|S \cap r| > \varepsilon |S|$ contain some point in N. Note that now N is not required to be a subset of S.

Such weak ε-nets turn out to be of little interest in our context for at least two reasons. For one, as we pointed out in the introduction, the range spaces $(\mathrm{R}^d, \mathcal{H}_d)$ are worth concentrating on since by the observation of Yao and Yao [YY85] many other geometric range spaces can be mapped to $(\mathrm{R}^d, \mathcal{H}_d)$. Typically this mapping is not invertible, and thus usually a weak ε-net from $(\mathrm{R}^d, \mathcal{H}_d)$ cannot be mapped back to the original range space of interest. This problem does not arise with "strong" ε-nets.

Moreover, in the case of $(\mathrm{R}^d, \mathcal{H}_d)$ a trivial weak ε-net of size $d+1$ exists for every finite set $S \subset \mathrm{R}^d$ and every ε: Simply take the vertices of a sufficiently large d-simplex that has all of S in its interior. Clearly any halfspace that contains a point of S also contains one of the vertices of that simplex.

We end this paper with a similar result about balls in R^d:

Theorem 3 *Let $d > 0$ be fixed and let $0 < \varepsilon \leq 1$. For every finite set $S \subset \mathrm{R}^d$ there exists a weak ε-net with respect to balls that has size at most $O(1/\varepsilon)$.*

Proof: This is an immediate consequence of a theorem of Wegner [W] which states that there is a constant c_d depending on d only so that if in a set B of balls in R^d among any k balls there are two that intersect, then there exists a set N of at most $c_d k$ points that has non-empty intersection with every ball in B. In our case let B be the collection of all balls that contain at least εn points of S and let k be $\lceil 1/\varepsilon \rceil + 1$. By the pigeon-hole principle some two out of any k such balls must have a point (of S) in common. $\square$

References

[ABFK91] N. Alon, I. Bárány, Z. Füredi, and D. Kleitman. Point selections and weak ε-nets for convex hulls. Manuscript, 1991.

[Aga90] P. K. Agarwal. Partitioning arrangements of lines I: An efficient deterministic algorithm. *Discrete & Computational Geometry*, 5:449–483, 1990.

[AK91] N. Alon and D. Kleitman. Piercing convex sets and the Hadwiger Debrunner (p,q)-problem. *Adv. Math.*, to appear. Also to appear in *Proc. 8. ACM Symposium on Computational Geometry*, 1992.

[BEHW89] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth. Classifying learnable geometric concepts with the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, 1989.

[CF90] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.

[Cha91] B. Chazelle. Cutting hyperplanes for divide-and-conquer. Tech. report CS-TR-335-91, Princeton University, 1991. Preliminary version: *Proc. 32. IEEE Symposium on Foundations of Computer Science*, October 1991.

[Cla87] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete & Computational Geometry*, 2:195–222, 1987.

[Cla89] K. Clarkson, unpublished proof outline, December 1989.

[CW89] B. Chazelle and E. Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete & Computational Geometry*, 4:467–490, 1989.

[ES86] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1:25–44, 1986.

[EW85] H. Edelsbrunner and E. Welzl. On the number of line separations of a finite set in the plane. *Journal of Combinatorial Theory Ser. A*, 38:15–29, 1985.

[HW87] D. Haussler and E. Welzl. ε-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[KPW91] J. Komlós, J. Pach, and G. Wöginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 1992. To appear.

[Mat90] J. Matoušek. Construction of ε-nets. *Discrete & Computational Geometry*, 5:427–448, 1990.

[Mat91a] J. Matoušek. Cutting hyperplane arrangements. *Discrete & Computational Geometry*, 6(5):385–406, 1991.

[Mat91b] J. Matoušek. Reporting points in halfspaces. In *Proc. 32. IEEE Symposium on Foundations of Computer Science*, pages 207–215, 1991.

[PW] J. Pach and G. Wöginger. Some new bounds for epsilon-nets. In *Proc. 6th ACM Symposium on Computational Geometry*, pages 10–15, 1990.

[VC71] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16:264–280, 1971.

[YY85] F. F. Yao and A. C. Yao. A general approach to geometric queries. In *Proc. 17. ACM Symposium on Theory of Computing*, pages 163–168, 1985.