Available online at www.sciencedirect.com

ScienceDirect

Computational Geometry 40 (2008) 84–91

Computational Geometry Theory and Applications

www.elsevier.com/locate/comgeo

Weak ϵ-nets have basis of size $O(1/\epsilon \log(1/\epsilon))$ in any dimension

Nabil H. Mustafa [a], Saurabh Ray [b,*]

[a] *Lahore University of Management Sciences, Pakistan*
[b] *Universitaet des Saarlandes, Saarbruecken, Germany*

Received 4 August 2006; received in revised form 4 January 2007; accepted 13 February 2007

Available online 4 April 2007

Communicated by E. Wezl

Abstract

Given a set P of n points in $\mathbb{R}^d$ and $\epsilon > 0$, we consider the problem of constructing weak ϵ-nets for P. We show the following: pick a random sample Q of size $O(1/\epsilon \log(1/\epsilon))$ from P. Then, with constant probability, a weak ϵ-net of P can be constructed from only the points of Q. This shows that weak ϵ-nets in $\mathbb{R}^d$ can be computed from a subset of P of size $O(1/\epsilon \log(1/\epsilon))$ with only the constant of proportionality depending on the dimension, unlike all previous work where the size of the subset had the dimension in the exponent of $1/\epsilon$. However, our final weak ϵ-nets still have a large size (with the dimension appearing in the exponent of $1/\epsilon$).

Keywords: Combinatorial geometry; Weak ϵ-nets; Hitting convex sets

1. Introduction

Given a set system $(X, \mathcal{F})$, where X is the base set, and $\mathcal{F}$ is a family of subsets of X, the general ϵ-net problem asks for a small subset X' of X such that for every set $S \in \mathcal{F}$ containing at least $\epsilon|X|$ elements, $X' \cap S \neq \emptyset$. In a celebrated result, Haussler and Welzl [5] showed that if the set system has finite VC-dimension, then picking a random sample from X of size $O(1/\epsilon \log(1/\epsilon))$ (constant dependent linearly on the VC-dimension of the set system) yields an ϵ-net with some constant probability. Subsequently the ϵ-net problem for systems of finite VC-dimension has been studied extensively [6].

Unfortunately, the existence of small ϵ-nets is no longer true for set systems of infinite VC-dimension. For example, it is easy to see that any ϵ-net with respect to convex ranges must have at least $(1-\epsilon)n$ points of P if P is in convex position. The concept of *weak* ϵ-nets with respect to *convex ranges* was introduced by Haussler and Welzl [5] in their seminal paper: the restriction that the points of ϵ-net be a subset of X is dropped. Weak ϵ-nets (w.r.t. convex ranges) have found several applications in discrete and combinatorial geometry (see Matousek's book for several examples [6]).

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. This is finite since Alon et al. [2] have shown that for any ϵ, d, there exist a weak ϵ-net of size independent of n. In

* Corresponding author.
E-mail addresses: nabil@lums.edu.pk (N.H. Mustafa), saurabh@cs.uni-sb.de (S. Ray).

0925-7721/$ – see front matter
doi:10.1016/j.comgeo.2007.02.006

particular, they proved that $w(d,\epsilon) \leqslant O(1/\epsilon^{d+1-\delta_d})$, where δ_d tends to zero with $d \to \infty$. This result was improved by Chazelle et al. [3] to $w(d,\epsilon) \leqslant O(1/\epsilon^d \operatorname{polylog}(1/\epsilon))$. 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 [7] gave an elegant algorithm that computes weak ϵ-nets in $\mathbb{R}^d$ of size $O(1/\epsilon^d \operatorname{polylog}(1/\epsilon))$. Their basic idea is the following: given the set P in $\mathbb{R}^d$, first compute a r-simplicial partition of P, r to be set later. Let S be the set formed by choosing an arbitrary point from each subset, and compute a set A (shown to be of size $O(r^{d^2})$) such that a centerpoint of every subset of S is present in A. The central claim is that if a convex set contains points from a large number of the sets of the partition, then it must contain the centerpoint of those points of S chosen from these intersected sets. Otherwise if the convex set intersects few sets of the partition, then Matousek and Wagner [7] recurse on the sets.

1.1. Our contributions

A long-standing open problem has been to show the existence of weak ϵ-nets in $\mathbb{R}^d$ with size $o(1/\epsilon^d)$. 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. [7] is that optimal weak ϵ-nets should have size $O(1/\epsilon \operatorname{polylog}(1/\epsilon))$ in $\mathbb{R}^d$ for every integer d. This conjecture and the following observation (which follows from Lemma 5.1) is the motivation for our work:

Observation 1.1. 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 $O(1/\epsilon)$ in any dimension, then there must exist $O(1/\epsilon)$ (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 $\Omega(1/\epsilon^d)$ input points*. In fact, a modification of [7] 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 $\Omega(1/\epsilon^d)$ input points are used.

Our contributions in this paper are threefold:

- We answer the above question in the affirmative, showing that for every point set P, there exists a set of $O(1/\epsilon \log(1/\epsilon))$ 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 $\Theta(1/\epsilon \log^{d^2}(1/\epsilon)$, 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 $O(1/\epsilon \log(1/\epsilon))$ points, one can construct a weak ϵ-net of size $O(1/\epsilon^3 \log^3(1/\epsilon))$. Similarly, we show that by random sampling $O(1/\epsilon \log(1/\epsilon))$ points in $\mathbb{R}^3$, and taking some function of them, one gets a weak ϵ-net of size $O(1/\epsilon^5 \log^5(1/\epsilon))$. For P in $\mathbb{R}^d$, take a random sample of size $O(1/\epsilon \log(1/\epsilon))$ (with only the constant depending on d). Then another product function of these sampled points yields an ϵ-net with size $O(1/\epsilon^{d^2})$.
- 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 5.1 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 $O(1/\epsilon^t \log^t(1/\epsilon))$. Therefore if one could show $t < d$, it improves the size of weak ϵ-nets.

1.2. Organization

We first present an elementary proof for the two-dimensional case in Section 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 4. The general construction for arbitrary dimensions is then presented in Section 5.

2. Preliminaries

We define a few concepts from discrete geometry for later use [6].

VC-dimension and ϵ-nets. (See [6].) Given a range space $(X, \mathcal{F})$, a set $X' \subseteq X$ is *shattered* if every subset of X' can be obtained by intersecting X' with a member of the family $\mathcal{F}$. The VC-dimension of $(X, \mathcal{F})$ is the size of the largest set that can be shattered. The ϵ-net theorem (Welzl and Haussler [5]) states that there exists an ϵ-net of size $O(d/\epsilon \log(1/\epsilon))$ for any range space with VC-dimension d.

Radon's theorem. *(See [6].) 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. *(See [4].) 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. *(See [1].) Given a set S of convex sets in $\mathbb{R}^d$ such that out of every $p \geqslant d+1$ set, there is a point common to $q \geqslant d+1$ 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 $|S|$).*

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 [5], it follows that a random sample of size $O(1/\epsilon \log(1/\epsilon))$ 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 3.1. *Let P be a set of n points in the plane, and let Q be an ϵ-net for the range space $\mathcal{R}_k$. Then, for* any *convex set $\mathcal{C}$ in the plane containing at least ϵn points of P, either* (a) *$\mathcal{C} \cap Q \neq \emptyset$, 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 $\mathcal{C}$ intersects the edge $\overline{q_i q_j}$ for all $1 \leqslant i < j \leqslant \lfloor k/2 \rfloor$.*

Proof. Assume $\mathcal{C} \cap Q = \emptyset$. 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 $\mathcal{C}' \subseteq C$ such that $\mathcal{C}' \cap P = \mathcal{C} \cap P$), 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 Fig. 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 Fig. 1(b).

Claim 3.1. *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 $\mathcal{C}$. Claim 3.1 immediately implies the following.

Claim 3.2. *The halfspace h_t, defined by the line incident to the edge $p_{i_t-1}p_{i_t}$, separates q_t (and all the other points of Q lying in $T(i_{t-1}, i_t)$) from $\mathcal{C}$.*

If the number of points found by our method is at most k (i.e., $l \leqslant k$), then take the intersection of the half-spaces h_t, for $t = 1, \ldots, l$. By Claim 3.2, each halfspace h_t separates all the points in $T(i_{t-1}, i_t)$ from $\mathcal{C}$. Thus all the points of Q are now separated by this intersection (see Fig. 1(a) for the separating halfplanes), and since each halfspace contains $\mathcal{C}$, 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 $\mathcal{C}$. This follows from the fact that for every point q_t, all the non-adjacent points and q_t lie in the same two half-spaces incident to edges $p_{i_t-1}p_{i_t}$ and $p_{i_{t+1}}p_{i_{t+1}+1}$, both of which are incident to $\mathcal{C}$. Therefore picking every alternate point yields the desired set. □

Set $k = 8$, and compute the ϵ-net for the range space $\mathcal{R}_8$. It follows from Lemma 3.1 that if a convex set $\mathcal{C}$ is not hit by the computed ϵ-net,then there exists a sequence of four points, say a, b, c, d, such that $\mathcal{C}$ contains the intersection of the two segments ac and bd. This immediately yields a way to construct weak ϵ-nets using (strong) ϵ-nets: 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 $O(1/\epsilon^4 \log^4(1/\epsilon))$. We now show that by a more careful argument, this can be reduced to $O(1/\epsilon^3 \log^3(1/\epsilon))$.

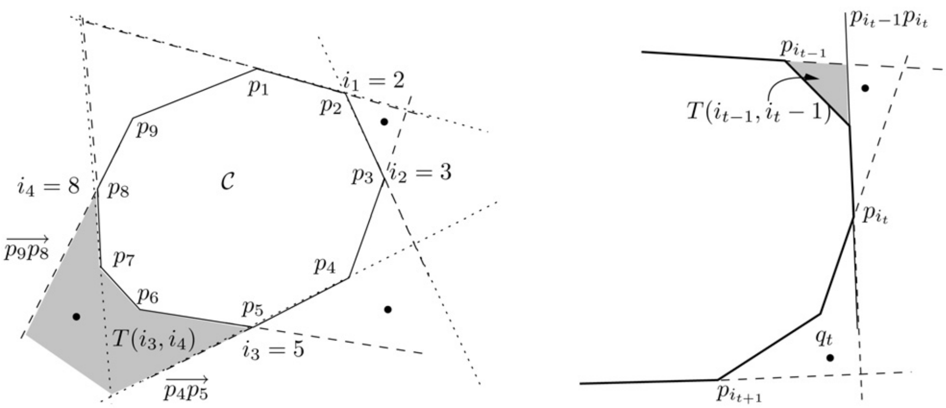

(a) (b)

Fig. 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 $\mathcal{C}$.

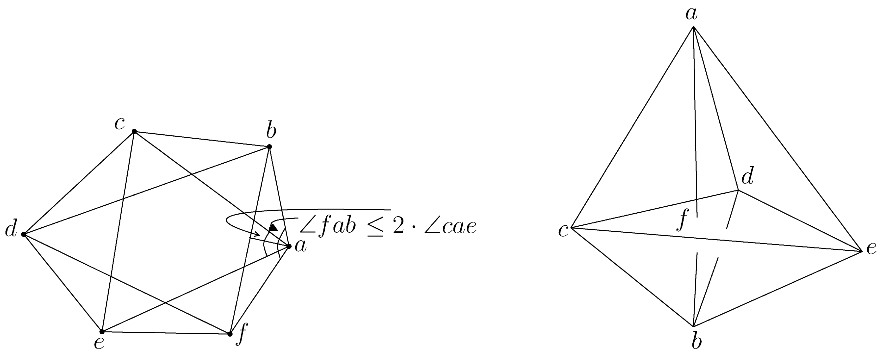

(a) (b)

Fig. 2. (a) The intersection of a bisector with a segment will lie inside $\mathcal{C}$, (b) If $\mathcal{C}$ intersects edges ac, ad and ae, then it must intersect af. Similarly for bf.

Theorem 3.1. *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' has size* $O(1/\epsilon^3 \log^3(1/\epsilon))$ *and is a weak ϵ-net for P.*

Proof. Fix a convex set $\mathcal{C}$ containing at least ϵn points of P. We may assume that $\mathcal{C}$ does not contain any point of Q. Then, from Lemma 3.1, there exists a sequence of six points in convex position, say a, b, c, d, e, f, of Q where the intersection point of every pair of (properly intersecting) segments spanning these points lies in $\mathcal{C}$.

The sum of the interior angles of the polygon defined by the six points is 4π. Form two triangles by taking alternate points, say $\triangle ace$ and $\triangle bdf$. The sum of the interior angles of the two triangles is 2π. By 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, $\angle fab$. 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 Fig. 2(a). By assumption, these two intersections are contained inside $\mathcal{C}$. Therefore, by convexity, the intersection of the bisector of $\angle fab$ with the segment fb lies inside $\mathcal{C}$. Since Q' contains all such intersections, $\mathcal{C}$ is hit by Q'. □

Remark. An alternate proof follows from the fact that given any point set P in $\mathbb{R}^2$, there exist 2 orthogonal lines which equipartition P [8].

4. Three dimensions

Lemma 4.1. *There exists a constant $f_d(t)$ for every $t \geqslant d + 1$ such that given a polytope $\mathcal{C}$ and a set of points Q in $\mathbb{R}^d$ such that $\mathcal{C} \cap Q = \emptyset$,* (i) *either the set Q can be separated from $\mathcal{C}$ by $f_d(t)$ hyperplanes or* (ii) *there exists $Q' \subseteq Q$ such that $|Q'| = t$ and the convex hull of every $d + 1$ points of Q' intersects $\mathcal{C}$.*

Proof. Assume, without loss of generality, that the origin lies in the interior of $\mathcal{C}$. For $\vec{q} \in Q$ define

$$S(\vec{q}) = \{\vec{a} \in \mathbb{R}^d \mid \vec{a} \cdot \vec{q} \geqslant 1, \vec{a} \cdot \vec{x} \leqslant 1 \ \forall x \in \mathcal{C}\},$$

where '·' denotes the inner product. First note that $S(\vec{q}) \neq \emptyset$ since $q \notin \mathcal{C}$. Second, $S(\vec{q})$ 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 $\mathcal{C}$). Since $\mathcal{C}$ contains the origin, $S(\vec{q})$ is also bounded and hence compact.

Since $\vec{0} \notin S(\vec{q})$, $\vec{a} \in S(\vec{q})$ implies that there is a hyperplane ($\vec{a} \cdot \vec{x} = 1$) which separates the point $\vec{q}$ from the $\mathcal{C}$. If there are $d + 1$ points $q_1, \ldots, q_{d+1}$ whose convex hull does not intersect $\mathcal{C}$, then these $d + 1$ points can be separated from $\mathcal{C}$ by a single hyperplane (separation theorem, [6]). This implies that the corresponding convex sets $S(\vec{q_1}), \ldots, S(\overrightarrow{q_{d+1}})$ have a common intersection.

Let $S = \{S(\vec{q}) | \vec{q} \in Q\}$ 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 $\mathcal{C}$, 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 $\mathcal{C}$ 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 4.2. *There exists a constant $g(t)$ for every $t \geqslant 5$ such that given a convex set $\mathcal{C}$ in $\mathbb{R}^3$ and set Q' of $g(t)$ points in $\mathbb{R}^3$ where the convex hull of every 4 points in Q' intersects $\mathcal{C}$, one can find $Q'' \subseteq Q'$ of size at least t such that the convex hull of every 3 points in Q'' intersects $\mathcal{C}$.*

Proof. Consider a hypergraph with the base set Q' and every 3-tuple of points in Q' as a hyperedge. Color a hyperedge 'red' if the convex hull of the corresponding 3 points intersects $\mathcal{C}$ and 'blue' otherwise. Then, by Ramsey's theorem for hypergraphs [4], there exists a constant $g(t)$ such that if $|Q'| \geqslant g(t)$, there exists a monochromatic clique, say Q'', of size t. A monochromatic 'blue' clique implies that there exists a set of t points such that $\mathcal{C}$ does not intersect the convex hull of any 3-tuple of these points. Take any 5 points of Q'', and partition their convex hull into two tetrahedra sharing a face. Since both these tetrahedra must intersect $\mathcal{C}$, their common face must also intersect $\mathcal{C}$, a contradiction. Therefore, the clique returned must be monochromatic 'red', implying the existence of a subset Q'' of size t such that the convex hull of all three points in Q'' intersects $\mathcal{C}$. □

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

Claim 4.1. *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 $\mathcal{C}$ intersects the convex hull of every 3-tuple of T, it intersects at least one edge (convex hull of a 2-tuple) spanned by the points in T.*

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

Lemma 4.3. *Given a convex set $\mathcal{C}$ in $\mathbb{R}^3$, there exists a constant $h(t)$ such that for any set Q'' of $h(t)$ points where the convex hull of every 3 points in Q'' intersects $\mathcal{C}$, one can find a subset $Q''' \subseteq Q''$ of size t such that the convex hull of every two points in Q''' intersects $\mathcal{C}$.*

Proof. Again consider a hypergraph with the base set Q'' and every 2-tuples of these points as a hyperedge. Color a hyperedge 'red' if the convex hull of the corresponding 2-tuple intersects $\mathcal{C}$ and 'blue' otherwise. Then again by Ramsey's theorem, there exists a positive integer $h(t)$ such that if $|Q''| \geqslant h(t)$, there exists a monochromatic clique of size t. We can assume (again by Ramsey's theorem) that if $t \geqslant k$ where k is a constant, then the points of the monochromatic clique have 5 points in convex position. From Claim 4.1, it follows that the convex hull of two of the points of these 5 points intersects $\mathcal{C}$, thereby implying that the color of the monochromatic clique cannot be 'blue' and hence the convex hull of every pair of points in the clique intersects $\mathcal{C}$. □

Lemma 4.4. *Given a set of points R in convex position in $\mathbb{R}^3$, $|R| \geqslant 5$, and a convex set $\mathcal{C}$ that intersects every edge spanned by the points in R, a Radon point of R is contained in $\mathcal{C}$.*

Proof. Take the Radon partition of any five points in R. See Fig. 2(b). Say the edge ab intersects the facet spanned by $\{c, d, e\}$. It is easy to see that if $\mathcal{C}$ intersects the edges ac, ad and ae, it must intersect the segment af. Similarly, if $\mathcal{C}$ intersects the edges bc, bd and be, it intersects the segment bf. By convexity, it must contain the intersection of the edge ab with $\triangle cde$. □

We come to our main theorem in this section:

Theorem 4.1. *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' is a weak ϵ-net for P of size $O(1/\epsilon^5 \log^5(1/\epsilon))$.*

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

For a large enough constant c (depending on $f_d(\cdot), g(\cdot), h(\cdot)$), by Lemmas 4.1, 4.2 and 4.3, there exists a set of at least five points such that $\mathcal{C}$ intersects every edge spanned by these points. Lemma 4.4 then implies that Q' is a weak ϵ-net. □

Remark. In [7], 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. Higher dimensions

Although the optimal weak ϵ-net can consist of any subset of $\mathbb{R}^d$, arguing similar to [7], we show that there is a discrete finite set of points in $\mathbb{R}^d$ from which an optimal weak ϵ-net can be chosen. Given P, this subset is constructed as follows: consider the set of all hyperplanes spanned by the points of P (each such hyperplane is defined by d points of P). Every d of these hyperplanes intersect in a point in $\mathbb{R}^d$. 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 5.1. *Let P be a set of n points in $\mathbb{R}^d$. Then the set $\Xi(P)$, of size $\mathrm{O}(n^{d^2})$, 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 [6]. 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. □

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

Theorem 5.1. *Let P be a set of n points in $\mathbb{R}^d$, and let Q be a random sample of size $\mathrm{O}(1/\epsilon \log(1/\epsilon))$ from P. With constant probability, $Q' = Q \cup \Xi(Q)$ is a weak ϵ-net for P.*

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

First, with constant probability, Q is an ϵ-net with respect to the range space (P, $\mathcal{R}_c$) for $c = f_d((d+1)^2)$, where $f_d(\cdot)$ is as in Lemma 4.1. Let $\mathcal{C}$ be any convex set containing at least ϵn points of P and assume $\mathcal{C} \cap Q = \emptyset$. Then $\mathcal{C}$ cannot be separated from Q by c hyperplanes, otherwise the intersection of the halfspaces containing $\mathcal{C}$ defined by these c hyperplanes has ϵn points and no point of Q, a contradiction to the fact that Q is an ϵ-net for (P, $\mathcal{R}_c$). Again assume, as in Theorem 4.1, that $\mathcal{C}$ is a full-dimensional polytope. By Lemma 4.1, there exist a set S of at least $(d+1)^2$ points of Q such that the convex hull of every $d+1$ of them intersects $\mathcal{C}$.

By Lemma 1 of [7], Q' contains a centerpoint, say q, of the set S. We claim that q is contained in $\mathcal{C}$. Otherwise, by the separation theorem, there exists a halfspace h^- containing q such that $h^- \cap \mathcal{C} = \emptyset$. By the centerpoint property, h^- contains at least $(d+1)^2/(d+1) = d+1$ points of S. The convex hull of these $d+1$ points lies in h^- and therefore does not intersect $\mathcal{C}$, a contradiction. □

Given a set Q, a *deep-point* is a point $q \in \mathbb{R}^d$ such that any halfspace containing q contains at least d points of Q. Let $c(Q)$ be the set of points in $\mathbb{R}^d$ such that a deep-point of every subset of Q of size at least $(d+1)^2$ is present in $c(Q)$. The proof above implies the following.

Corollary 5.1. *If $c(Q)$ has size $\mathrm{O}(m^t)$ for any set Q of size m, one can construct a weak ϵ-net for any point set of size* $\mathrm{O}(1/\epsilon^t \log^t(1/\epsilon))$.

Acknowledgements

The second author would like to thank Raimund Seidel and Hansraj Tiwary for insightful discussions.

References

[1] N. Alon, D. Kleitman, Piercing convex sets and the Hadwiger Debrunner $(p;q)$-problem, Adv. Math. 96 (1) (1992) 103–112.
[2] N. Alon, I. Bárány, Z. Füredi, D.J. Kleitman, Point selections and weak ε-nets for convex hulls, Combin. Probab. Comput. 1 (1992) 189–200.
[3] B. Chazelle, H. Edelsbrunner, M. Grigni, L. Guibas, M. Sharir, E. Welzl, Improved bounds on weak ε-nets for convex sets, in: STOC '93: Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, 1993, pp. 495–504.
[4] R. Diestel, Graph Theory, Springer-Verlag, New York, 2000.
[5] D. Haussler, E. Welzl, ε-nets and simplex range queries, Discrete Comput. Geom. 2 (1987) 127–151.
[6] J. Matousek, Lectures in Discrete Geometry, Springer-Verlag, New York, 2000.
[7] J. Matousek, U. Wagner, New constructions of weak ε-nets, Discrete Comput. Geom. 32 (2) (2004) 195–206.
[8] E.A. Ramos, Equipartition of mass distributions by hyperplanes, Discrete Comput. Geom. 15 (2) (1996) 147–167.