Discrete Comput Geom 7:163–173 (1992)

Discrete & Computational
Geometry

Almost Tight Bounds for ε-Nets

János Komlós,[1] János Pach,[2,3] and Gerhard Woeginger[4]

[1] Department of Mathematics, Rutgers University,
New Brunswick, NJ 08903, USA

[2] Mathematical Institute, Hungarian Academy of Sciences,
Pf. 127, H-1364 Budapest, Hungary

[3] Courant Institute of Mathematical Sciences, New York University,
New York, NY 10012, USA

[4] Institut für Mathematik, Technische Universität Graz,
Kopernikusgasse 24, A-8010 Graz, Austria

Abstract. Given any natural number d, $0 < \varepsilon < 1$, let $f_d(\varepsilon)$ denote the smallest integer f such that every range space of Vapnik–Chervonenkis dimension d has an ε-net of size at most f. We solve a problem of Haussler and Welzl by showing that if $d \geq 2$, then

$$d - 2 + \frac{2}{d+2} \leq \lim_{\varepsilon \to 0} \frac{f_d(\varepsilon)}{(1/\varepsilon)\log(1/\varepsilon)} \leq d.$$

Further, we prove that $f_1(\varepsilon) = \max(2, \lceil 1/\varepsilon \rceil - 1)$, and similar bounds are established for some special classes of range spaces of Vapnik–Chervonenkis dimension three.

1. Introduction

ε-nets were introduced by Haussler and Welzl [8]. The concept proved to be useful in many fields of discrete and computational geometry and in learnability theory, see [1], [2], [4], [6], and [9]. In this paper we show that the upper bounds for $f_d(\varepsilon)$ given in [8] and [2] are optimal up to a constant factor.

The following terminology is taken from [8]: A *range space* S is a pair (X, R), where X is a set and R is a family of subsets of X. The members of X are called *points* or *elements*, members of R are called *ranges*. The range space is *finite* if X is finite. A subset $A \subset X$ is said to be *shattered* by R, if every subset of A can be obtained by intersecting A with some range in R. The *Vapnik–Chervonenkis dimension* of the range space (or *VC-dimension*, for short) is the cardinality of the largest shattered subset of X. If arbitrarily large sets can be shattered, then the VC-dimension is *infinite*.

For some real ε, $0 < \varepsilon < 1$, a subset N of a finite set $T \subset X$ is called an *ε-net for T with respect to R*, if N contains at least one point from each range $r \in R$ with $|r \cap T| > \varepsilon|T|$.

Very often, especially in discrete geometry, the actual range space $S' = (X', R')$ is finite but is obtained as the restriction to X' of an infinite range space (X, R), where $X' \subset X$ and $R' = \{X' \cap r; r \in R\}$ (see the example below). In most cases, only the range space (X, R) matters, and the size or the actual choice of X' does not. The following more general definition (the original Vapnik–Chervonenkis set-up) is an alternative to restricting an infinite X to a finite X'.

Given a possibly infinite range space $S = (X, R)$, and a probability measure μ on X, a set $N \subset X$ is called an ε-net if N intersects all ranges $r \in R$, $\mu(r) > \varepsilon$. Typically μ is the uniform measure on a finite set X'. We simply use the word "restriction" in any case.

In [8] the following proposition is proved by using the probabilistic method of Vapnik and Chervonenkis [13].

Proposition. *Let $S = (X, R)$ be a finite range space of VC-dimension d, and let $0 < \varepsilon < 1$. Then there exists an ε-net for X with respect to R of size at most $\lceil (8d/\varepsilon) \cdot \log(8d/\varepsilon) \rceil$. (Note that the bound is independent of the size of X.)*

Example. Let X be a finite set of n points in the Euclidean plane, and let R be the restriction to X of the family of all open half-planes. The VC-dimension of such a range space is at most three. Indeed, consider any set of four points in the plane. If one of the points is contained in the convex hull of the other points, the inner point cannot be cut off alone. If all four points are extreme, two diametrical points cannot be cut off.

For given ε between 0 and 1, an ε-net for this range space must contain a point in each half-plane that cuts off more than εn points from X. By the above theorem of Haussler and Welzl, there exists an ε-net of size at most $\lceil (24/\varepsilon) \log(24/\varepsilon) \rceil$. In Section 5 of this paper we show that this bound can be improved to $\lceil 2/\varepsilon \rceil - 1$.

Furthermore, it was shown by Haussler and Welzl [8] that $f_d(\varepsilon) \geq d/(2\varepsilon) - 1$. They finished their paper by asking where the function $f_d(\varepsilon)$ actually lies between the two bounds $\Omega(d/\varepsilon)$ and $O((d/\varepsilon) \log(d/\varepsilon))$. As a first result, Blumer *et al.* [2] improved the upper bound to $O((d/\varepsilon) \log(1/\varepsilon))$. Then in some special applications of ε-nets to computational geometry, namely subdividing the space into simplices with respect to a given set of hyperplanes, various authors [1], [3], [9] managed to do away with the $\log(1/\varepsilon)$-factor for some special range spaces, building on the results of Clarkson [5].

In 1989 Matoušek *et al.* [10] succeeded in establishing the $O(1/\varepsilon)$ bound on the size of ε-nets for *half-space*, *disk*, and *pseudodisk range spaces*. That is, if we take R to be the set of all open half-spaces *in three-space*, or all disks or pseudodisks *in two-space*, there will be an $O(1/\varepsilon)$-size ε-net for any restriction of the range space. In Section 6 we present the same bound for ε-nets of range spaces, where R contains all translates of some fixed simple polygon in the plane. No similar bounds are known for Euclidean spaces of dimensions higher than three.

We note that in none of the known examples is the extra factor $\log(1/\varepsilon)$ necessary. In Section 2 we show that in general it is. In other words, we show that $O((d/\varepsilon)\log(1/\varepsilon))$, the upper bound on $f_d(\varepsilon)$ given by Blumer *et al.* [2], is optimal up to a constant factor for all $d \geq 2$. Our proof method is probabilistic and hence not constructive. However, considering results on discrepancy, it seems likely that the *triangle range space* gives a simple geometric example where the $\log(1/\varepsilon)$-factor is necessary. (In a triangle range space, X is a set of points in the unit square and R is the restriction to X of the family of all triangles.) Section 3 contains a slight improvement of the upper bound of [2]. For $d = 1$, we determine $f_1(\varepsilon)$ explicitly in Section 4. Section 5 deals with ε-nets for half-planes and Section 6 gives a similar result for the translates of fixed simple polygons.

2. A Lower Bound for Dimensions ≥ 2

In this and the next section we show, using probabilistic methods, that, for any given $d \geq 2$, $f_d(\varepsilon)$ is proportional to $(1/\varepsilon)\log(1/\varepsilon)$, where the constant involved is between $d - 2$ and d provided that ε is small enough.

More precisely, we prove

$$d - 2 + \frac{2}{d+2} \leq \liminf_{\varepsilon \to 0} \frac{f_d(\varepsilon)}{(1/\varepsilon)\log(1/\varepsilon)} \leq \limsup_{\varepsilon \to 0} \frac{f_d(\varepsilon)}{(1/\varepsilon)\log(1/\varepsilon)} \leq d.$$

We do not know if the limit above exists, so the statement in the abstract is somewhat sloppy.

Theorem 2.1. *Given any integer $d \geq 2$ and any real $\gamma < 2/(d+2)$, there exists an $\varepsilon_0(d, \gamma) > 0$ such that for every $\varepsilon \leq \varepsilon_0(d, \gamma)$ we can construct a finite range space $S = (X, R)$ of VC-dimension $\leq d$ which does not have an ε-net of size smaller than $(d - 2 + \gamma)(1/\varepsilon)\log(1/\varepsilon)$. In other words,*

$$\liminf_{\varepsilon \to 0} \frac{f_d(\varepsilon)}{(1/\varepsilon)\log(1/\varepsilon)} \geq d - 2 + \frac{2}{d+2}.$$

Proof. In the following, certain parameters (n, r, t) are supposed to be large integers, and we disregard the roundoff errors.

Let γ' be any constant between γ and $2/(d+2)$. Given a sufficiently small ε, let $n = K(1/\varepsilon)\log(1/\varepsilon)$, where the constant K depending only on d, γ, and γ' but *not* on ε, will be determined later. Let $r = \varepsilon n$, $p = \varepsilon^{1-d-\gamma'} \Big/ \binom{n}{r}$, and let $t = (d - 2 + \gamma)(1/\varepsilon)\log(1/\varepsilon)$. Let X be an n-element underlying set. The range set

R will consist of randomly selected r-element subsets of X, where each r-set has the same chance p to be selected and the selections are done independently.

We show that if the number n is sufficiently large (that is, if ε is sufficiently small), then the following two claims hold with a large probability:

(i) *The range space S has VC-dimension $\leq d$.*
(ii) *S does not allow an ε-net of size at most t.*

Proof of (*i*). The probability that the VC-dimension of S exceeds d is

$$\leq \binom{n}{d+1} \Big/ \text{Prob[a fixed } (d+1)\text{-element set is shattered by some } 2^{d+1} \text{ ranges]}$$

$$= \binom{n}{d+1} \prod_{i=0}^{d+1} \left[1 - (1-p)^{\binom{n-d-1}{r-d-1+i}}\right]^{\binom{d+1}{i}}$$

$$\leq \binom{n}{d+1} \prod_{i=0}^{1} \left[1 - (1-p)^{\binom{n-d-1}{r-d-1+i}}\right]^{\binom{d+1}{i}}$$

$$\leq \binom{n}{d+1}\left[\binom{n-d-1}{r-d-1}p\right]\left[\binom{n-d-1}{r-d}p\right]^{d+1} \leq n^{d+1}\left[\binom{n}{r}p\right]^{d+2} \varepsilon^{(d+1)^2}.$$

(Here we used the inequality

$$\binom{n-k}{r-k} \leq \binom{n}{r}\left(\frac{r}{n}\right)^k.)$$

Plugging in the selected expressions for n and p, we get the upper bound $(K \log 1/\varepsilon)^{d+1} \varepsilon^{2-(d+2)\gamma'}$, which goes to 0 as $\varepsilon \to 0$.

Proof of (*ii*). The probability that there exists an ε-net of size t for X is

$$\leq \binom{n}{t}(1-p)^{\binom{n-t}{r}} \leq \binom{n}{t} e^{-p\binom{n-t}{r}}.$$

We estimate $\binom{n}{t}$ from above by $(en/t)^t$, and $\binom{n-t}{r}$ from below by $\binom{n}{r}(1 - r/(n-t+1))^t$. Using the inequality $1 - ax > e^{-bx}$ for $b > a$, $0 < x < 1/a - 1/b$, we get the upper bound

$$\left(\frac{en}{t}\right)^t \exp\{-\varepsilon^{1-d-\gamma'} e^{-(K/(K-d))\varepsilon t}\} = \left(\frac{en}{t}\right)^t \exp\{-\varepsilon^{1-d-\gamma'+(K/(K-d))(d-2+\gamma)}\}.$$

This will go to 0 if the exponent $1 - d - \gamma' + (K/(K-d))(d-2+\gamma) < -1$, which holds if K was chosen large enough. □

3. An Upper Bound for Dimensions ≥ 2

In this section we give a better constant value in the upper bound of Blumer *et al.* [2]. Just as in [7] and [2], we simply adapt the original proof of Vapnik and Chervonenkis to the covering problem.

Theorem 3.1. *Let $S = (X, R)$ be an arbitrary range space of VC-dimension d, and let μ be an arbitrary probability measure on X. (We assume that the ranges are all μ-measurable.) If $\varepsilon > 0$ is sufficiently small in terms of d, then there exists an ε-net of size at most $t = (d/\varepsilon)[\log(1/\varepsilon) + 2\log\log(1/\varepsilon) + 3]$, that is, a set $N \subset X$ of size $\leq t$ such that N intersects every $r \in R$ with $\mu(r) \geq \varepsilon$.*

In fact, a random set of size t is an ε-net with a large probability (roughly $1 - e^{-d}$). Consequently,

$$\limsup_{\varepsilon \to 0} \frac{f_d(\varepsilon)}{(1/\varepsilon)\log(1/\varepsilon)} \leq d.$$

Just as in the previous section, we were somewhat sloppy in stating the theorem, for the value of t above is not an integer.

Proof. We use the main lemma of Vapnik and Chervonenkis [13], independently discovered by Sauer [12].

Given a subset $Y \subset X$, we write $R_{|Y}$ for the restriction of R to Y. Define

$$g(R, k) = \max_{|Y| = k} |R_{|Y}|.$$

Theorem (Vapnik–Chervonenkis, Sauer). *For any range space $S(X, R)$ of dimension d,*

$$g(R, k) \leq \sum_{i=0}^{d} \binom{k}{i}.$$

Now let us select with possible repetition t random points from the universe X, where the selections are done with respect to the measure μ. We get a random sample $x \in X^t$. As the bits are picked independently, x is selected with respect to the measure μ^t. Write N for the set of elements in x.

Theorem 3.2. *Let $T > t$ be any integer. Then*

$$\text{Prob}(N \textit{ does not cover } R) \leq 2g(R, T)\left(1 - \frac{t}{T}\right)^{(T-t)\varepsilon - 1}.$$

Choosing $T = (d/\varepsilon)(\log(1/\varepsilon))^2$, we get from this after some routine calculations that Prob(N does not cover R) < 1, which implies the upper bound claimed in Theorem 3.1. Thus, we only prove the last theorem.

Having picked the string x of length t, let us keep on choosing other $T - t$ elements. Call the new string y, and let $z = (x, y)$. For a specific $r \in R$, we write $I(r, x)$ for the number of bits in x that belong to r, counting with multiplicity.

We want to estimate the probability

$$\text{Prob}(N \text{ does not cover } R) = P(\exists r \in R;\ I(r, x) = 0).$$

The following inequality is an easy consequence of the independence of x and y, but the reader may want to reflect on this, since this together with the conditioning that follows are the heart of the proof:

$$\text{Prob}(\exists r \in R;\ I(r, x) = 0) \le \frac{\text{Prob}(\exists r \in R; I(r, x) = 0, I(r, y) \ge m)}{\min_{r \in R} \text{Prob}(I(r, y) \ge m)}$$

Choosing m to be the median of $I(r, y)$, we get

$$\text{Prob}(\exists r \in R;\ I(r, x) = 0) \le 2\ \text{Prob}(\exists r \in R;\ I(r, x) = 0,\ I(r, y) \ge m).$$

For a sample $z = (x, y)$, we write $\{x, y\}$ for the multiset of sample elements in (x, y). For a fixed r, the conditional probability for given $\{x, y\}$ is (writing $s = |I(r, z)|$)

$$\begin{aligned}\text{Prob}(I(r, x) = 0, I(r, y) \ge m \mid \{x, y\}) &\le \chi[I(r, z) \ge m]\binom{T - t}{s}\Big/\binom{T}{s} \\ &\le \chi[I(r, z) \ge m]\left(1 - \frac{t}{T}\right)^s \le \left(1 - \frac{t}{T}\right)^m.\end{aligned}$$

Since given $\{x, y\}$, there are at most $g(R, T)$ different intersections of the r's with the multiset $\{x, y\}$, and we get

$$\text{Prob}(\exists r \in R;\ I(r, x) = 0,\ I(r, y) \ge m \mid \{x, y\}) \le g(R, T)\left(1 - \frac{t}{T}\right)^m.$$

Taking total expectations, and using the known fact that the median of a binomial distribution is within 1 of the mean, $m \ge (T - t)\varepsilon - 1$, we get the theorem. □

4. Exact Bounds for Dimension One

In this section we determine the exact value of $f_1(\varepsilon)$, that is, the smallest size of an ε-net guaranteed in every range space of VC-dimension one.

Theorem 4.1. *For every finite range space $S = (X, R)$ of VC-dimension one and for any real number ε, $0 < \varepsilon < 1$, there exists an ε-net of size at most* $\max(2, \lceil 1/\varepsilon \rceil - 1)$.

Proof. As the range space S is one dimensional, every pair of distinct elements $x, y \in X$ satisfies at least one of the following four conditions:

(i) There is no range $r \in R$ such that $r \cap \{x, y\} = \emptyset$ holds.
(ii) There is no range $r \in R$ such that $r \cap \{x, y\} = \{x\}$ holds.
(iii) There is no range $r \in R$ such that $r \cap \{x, y\} = \{y\}$ holds.
(iv) There is no range $r \in R$ such that $r \supseteq \{x, y\}$ holds.

If condition (i) holds for some pair $\{x, y\}$, we choose these two elements to form the ε-net and the theorem is true. Otherwise, we define a partial order "$\prec$" as follows:

If (ii) is valid for some pair $x, y \in X$, then let $x \prec y$.
If (iii) is valid for some pair $x, y \in X$, then let $y \prec x$.

If $x \prec y$ and $y \prec x$, then we say that x is equivalent to y ($x \sim y$). That is, x and y are indistinguishable: every $r \in R$ contains either both or none of them. It is clear that "$\prec$" defines a partial order $\mathscr{P}$ on the equivalence classes. (We disregard those elements that do not belong to any range.)

We claim that $\mathscr{P}$ has at most $\lceil 1/\varepsilon \rceil - 1$ maximal elements. Indeed, it is easy to see that each connected component of $\mathscr{P}$ contains exactly one maximal element, furthermore, each connected component $\mathscr{C}$ of $\mathscr{P}$ has more than εn elements. (Consider some element x in $\mathscr{C}$. x appears in at least one range. This range is entirely contained in $\mathscr{C}$ and its cardinality is larger than εn.)

Moreover, every range $r \in R$ contains a maximal element of $\mathscr{P}$ (trivial). Hence, the at most $\lceil 1/\varepsilon \rceil - 1$ maximal elements form an ε-net. □

Theorem 4.2. *For every real* ε, $0 < \varepsilon < 1$, *there exists a one-dimensional range space that does not allow an* ε*-net of size less than* $\max(2, \lceil 1/\varepsilon \rceil - 1)$.

Proof. (Case 1) If $\varepsilon \geq \frac{1}{2}$, then $\max(2, \lceil 1/\varepsilon \rceil - 1)$ is equal 2. To show that this bound is tight, choose an integer n large enough to satisfy $(n-1)/n > \varepsilon$. Consider the range space $S = (X, R)$ with $X = \{1, 2, \ldots, n\}$ and $R = \{X - \{k\};\ 1 \leq k \leq n\}$. Obviously, S is of VC-dimension one, as for no pair $x, y \in X$ and for no range $r \in R$ is $\{x, y\} \cap r = \emptyset$ satisfied. On the other hand, an ε-net for S has to contain at least two elements.

(Case 2) If $\varepsilon < \frac{1}{2}$, then $\max(2, \lceil 1/\varepsilon \rceil - 1)$ is equal to $\lceil 1/\varepsilon \rceil - 1$. The claim follows from the following more general observation.

For any d and for any $0 < \varepsilon < 1$, it is easy to construct a range space of VC-dimension d which does not permit an ε-net of size smaller than $\lceil d/\varepsilon \rceil - d$. Indeed, choose an integer n such that $\varepsilon n/d$ is not an integer and let X be a set of cardinality n. Partition X into sets $(X_0, X_1, \ldots, X_k)$, such that $|X_0| < \lceil \varepsilon n/d \rceil$ and such that $|X_i| = \lceil \varepsilon n/d \rceil$, for $i > 0$ (X_0 may be empty). An easy calculation shows that k equals $\lceil d/\varepsilon \rceil - 1$, for n large enough.

Let R be the family of all subsets that can be obtained as the union of d distinct parts X_i in the above partition. Obviously, the range space $S = (X, R)$ is of VC-dimension d. Now consider an ε-net of S: If it does not contain more than

$(k-d)$ elements of X, then there exist d of the X_i's with $X_i \cap Net = \emptyset$. The union of these d sets is a range in R. The claim follows. □

5. ε-Nets for Half-planes

In this section we estimate the size of the smallest ε-net in a half-plane range space (see the example in Section 1).

Theorem 5.1. *Let X be a finite point set in the plane, $n = |X|$, let R be the restriction to X of the family of all open half-planes, and let $0 < \varepsilon < 1$. Then there exists an ε-net for X with respect to R of size at most $\lceil 2/\varepsilon \rceil - 1$.*

Proof. (Case 1) If $\frac{2}{3} \le \varepsilon < 1$ holds, then $\lceil 2/\varepsilon \rceil - 1 = 2$. In this case we use the following property of the so-called *center* c of X: every open half-plane covering more than $2n/3$ points of X contains c. (A proof for the existence of such a point c can be found, e.g., in [7].)

Let $\triangle xyz$ be a triangle in X containing c but no element of X in its interior. Let us partition the elements of X into three classes according to which side of $\triangle xyz$ is met by the segments connecting them with c. (Note that if $\overline{vc}$ passes through a vertex of $\triangle xyz$ for some $v \in X$, then v belongs to two classes.) Obviously, at least one of the three classes, say the class corresponding to the side $\overline{xy}$, has at least $n/3 + 1$ points. We choose $\{x, y\}$ to be our net. (Every half-plane *not containing* c, contains at most $2n/3$ points. Every half-plane *containing* c but neither x nor y avoids the $n/3$ points behind $\overline{xy}$.)

(Case 2) $0 < \varepsilon < \frac{2}{3}$ holds. Consider an ε-net N such that (i) all points of N lie on the convex hull of X and (ii) N is minimal, i.e., if we remove any point from N, it is not an ε-net any more. Let $p_0, p_2, \ldots, p_{k-1}$ be the points of N listed in clockwise order on the hull. For $0 \le i \le k-1$, let $L(i)$ denote the line through p_{i-1} and p_{i+1}, let $h(i)$ be the open half-plane bounded by $L(i)$ that contains p_i and let $|h(i)|$ be the number of points in $h(i) \cap X$. If $k = 3$, we already have an ε-net of size $3 \le \lceil 2/\varepsilon \rceil - 1$ and so we assume $k \ge 4$ in the rest of the proof.

We observe that $|h(i)| > \varepsilon n$ holds for all i, $0 \le i \le k-1$ (otherwise we could remove p_i from the net). Furthermore, we claim that every point of X lies in at most two of the $h(i)$'s: if $p \in X$ lies in three different $h(i)$'s, it must lie in some $h(i) \cap h(j)$, where $i > j + 1$, but in this case, the points p_{i-1} and p_{j+1} are not extreme, a contradiction. Thus we get that the $h(i)$'s altogether contain each point at most twice and therefore

$$2|X| \ge \sum_{i=1}^{k} |h(i)| > \sum_{i=1}^{k} \varepsilon |X| = k\varepsilon |X|$$

holds. This directly implies $k \le \lceil 2/\varepsilon \rceil - 1$. □

This result is close to being optimal, as the following theorem gives a lower bound that deviates from the upper bound by at most one.

Theorem 5.2. *For any ε, $0 < \varepsilon < 1$, there exists a range space (X, R), where X is a finite point set in the plane, and R is the restriction to X of the family of all open half-planes, such that every ε-net for (X, R) contains at least $2\lceil 1/\varepsilon \rceil - 2$ elements.*

Proof. We explicitly construct such a range space. The pointset X is defined as follows: $\lfloor \varepsilon n \rfloor + 2$ points are placed on a piece of the parabola $y = x^2$, $-1 < x < +1$. Then every single point p in this group can be cut off by a line $l(p)$ that lies below the other $\lfloor \varepsilon n \rfloor + 1$ points in this group. We place $n/(\lfloor \varepsilon n \rfloor + 2)$ of these groups on the parabola $y = -x^2$, $-\infty < x < +\infty$, in such a way that each $l(p)$ lies below the $\lfloor \varepsilon n \rfloor + 1$ other points in the same group as p, but above all the other points.

Obviously, an ε-net for (X, R) must contain at least two points from each group (otherwise, $l(p)$ would cut off $\lfloor \varepsilon n \rfloor + 1$ points not in the net). If n is large enough, there are $\lceil 1/\varepsilon \rceil - 1$ groups. □

6. ε-Nets for the Translates of a Simple Polygon

Consider some fixed closed simple polygon P in the plane. In this section we deal with ε-nets for range spaces $S = (X, R_P)$ of the following form: X is a finite set of points in the plane and R_P is the restriction to X of the family of all translates P' of P. Obviously, for *nonconvex* polygons P, the VC-dimension of such polygon range spaces can become arbitrarily large, depending on the particular shape of P. For *convex* polygons P, we can prove the following lemma. The reader will observe that the proof neither makes use of the boundedness of the polygon nor of its piecewise straight boundary, but only of its convexity. Hence, the result holds for arbitrary convex regions in the plane.

Lemma 6.1. *For any convex polygon P and for any set X of points in the plane, the VC-dimension of (X, R_P) is at most three.*

Proof. We have to show that no set $\{a, b, c, d\}$ of four points in the plane can be shattered by the translates of P. If one of the points is contained in the convex hull of other points, these other points cannot be covered by P without covering the inner point at the same time. Hence, we may assume that the four points are in general position and form a convex quadrangle $\Box abcd$. We will prove that one of the diagonal pairs ac or bd cannot be cut off by P.

To do this, we look at the problem from a different point of view. Instead of keeping the four points fixed and covering them by translates of P, we now fix the polygon P and translate $\Box abcd$. If both diagonals could be covered by P, then there must exist two copies $\Box a'b'c'd'$ and $\Box a''b''c''d''$ such that a', c', b'', d'' lie inside of P and a'', c'', b', d' lie outside. For x, y, two consecutive points in $\Box a'b'c'd'$, let $h^+(x, y)$ denote the closed half-plane bounded by the line through x and y and containing $\Box a'b'c'd'$ and let $h^-(x, y)$ denote the complement of $h^+(x, y)$. Now the following holds:

b'', d'' cannot lie in $h^-(a', d') \cap h^-(d', c')$.
b'', d'' cannot lie in $h^-(a', b') \cap h^-(b', c')$.
a'' cannot lie in $h^+(d', a') \cap h^+(a', b')$.
c'' cannot lie in $h^+(d', c') \cap h^+(c', b')$.

(If b'' lay in $h^-(a', d') \cap h^-(d', c')$, then d' would lie in the interior of $\triangle a'c'b''$. As a', c', b'' lie inside P and P is convex, d' must be inside P, too. This is a contradiction. The other cases are settled by analogous arguments.) Now it is obvious that there is no possible position for $\square a''b''c''d''$. □

Theorem 6.2. *For any simple closed polygon P in the plane, there exists a constant $c(P)$ with the following property. For any polygon range space $S = (X, R_P)$ and for any real ε, $0 < \varepsilon < 1$, there exists an ε-net of size at most $c(P)/\varepsilon$.*

Proof. We use arguments similar to those in [11]. First, we observe that it is sufficient to prove the theorem for triangles, because the interior of any simple closed polygon can be decomposed into a set of triangles $\tau_1, \tau_2, \ldots, \tau_k$. If our polygon P covers $\geq \varepsilon n$ points of X, then at least one of these triangles covers $\geq (\varepsilon/k)n$ points. Hence, if we can find, for each triangle $\tau_1, \tau_2, \ldots, \tau_k$, (ε/k)-nets $N_1, N_2, \ldots, N_k$ of sizes at most $c(\tau_1)/(\varepsilon/k), \ldots, c(\tau_k)/(\varepsilon/k)$, then the set $N = N_1 \cup \cdots \cup N_k$ will meet the requirements with $c(P) = k(c(\tau_1) + \cdots + c(\tau_k))$.

Further, if P is a triangle we can do the following. We divide the plane into squares by putting a grid on it. Let us choose the size of the grid (i.e., the length of the squaresides) so small that the diameter of each square is smaller than the width (i.e., the smallest altitude) of the triangle P. This ensures that any square intersects at most two sides of a translate of P.

Now for any translate of P containing at least εn elements of X, there exists a square such that this translate covers at least $(\varepsilon/t)n$ elements of X within this particular square, where t is the number of squares intersected by the translate of P. Clearly, $(\text{diam}(P)/\text{gridsize} + 1)^2$ is an upper bound for t. Now our goal is to find, within each square Q_i that contains $n_i \geq (\varepsilon/t)n$ points of X, an $(\varepsilon n/(tn_i))$-net M_i of size at most $\tilde{c}(P)/(\varepsilon n/(tn_i))$ for translates of P. Then $M = M_1 \cup M_2 \cup \cdots$ forms an ε-net for $S = (X, R_P)$ of size at most

$$\sum_i \frac{\tilde{c}(P)tn_i}{\varepsilon n} = \frac{\tilde{c}(P)t}{\varepsilon} \frac{\sum n_i}{n} \leq \frac{c(P)t}{\varepsilon}$$

and by setting $c(P) := t\tilde{c}(P)$ our proof is complete.

Thus, it is sufficient to prove the statement for point sets within some square Q_i. As any square Q_i is intersected by at most two sides of a translate of P, we are in a similar situation as in the proof for the half-plane range spaces in Section 4. We define a point x in $X \cap Q_i$ to be *P-extreme* if it can be separated from the other points in $X \cap Q_i$ by a translate of P, i.e., x lies inside the translate and $X \cap Q_i - \{x\}$ lies outside. The P-extreme points form a so-called *P-convex hull* and now we can apply an argument analogous to that in Theorem 5.1. □

References

1. P. K. Agarwal, A deterministic algorithm for partitioning arrangements of lines and its applications, *Discrete & Computational Geometry* **5**, 1990.
2. A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth, Classifying learnable geometric concepts with the Vapnik–Chervonenkis dimension, *Journal of the Association for Computing Machinery* **36**, 1989.
3. B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Proceedings of the 29th IEEE Foundations of Computer Science*, 1988, 539–549.
4. B. Chazelle and E. Welzl, Quasi-optimal range searching and VC-dimensions, *Discrete & Computational Geometry* **4**, 1989, 467–490.
5. K. Clarkson, New applications of random sampling in computational geometry, *Discrete & Computational Geometry* **2**, 1987, 195–222.
6. K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and surfaces, *Discrete & Computational Geometry* **5**, 1990, 99–160.
7. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, EATCS Monographs in Theoretical Computer Science, Springer-Verlag, Berlin, 1987.
8. D. Haussler and E. Welzl, ε-Nets and simplex range queries, *Discrete & Computational Geometry* **2**, 1987, 127–151.
9. J. Matoušek, Construction of ε-nets, *Proceedings of the 5th Annual Symposium on Computational Geometry*, Saarbrücken, 1989, 1–10.
10. J. Matoušek, R. Seidel, and E. Welzl, How to net a lot with little: small ε-nets for disks and halfspaces, *Proceedings of the 6th Annual Symposium on Computational Geometry*, Berkeley, 1990.
11. J. Pach, Covering the plane with convex polygons, *Discrete & Computational Geometry* **1**, 1986, 73–81.
12. N. Sauer, On the density of families of sets, *Journal of Combinatorial Theory, Series A* **13**, 1972, 145–147.
13. V. N. Vapnik and A. Ya. Chervonenkis, On the uniform convergence of events to their probabilities, *Theory of Probability and Its Applications* **16**, 1971, 264–280.

Received December 1, 1989, and in revised form July 3, 1990.