Epsilon-Nets and Simplex Range Queries

David Haussler and Emo Welzl*

Department of Mathematics and Computer Science, University of Denver, Denver, Colorado 80208, USA.
*Institutes for Information Processing, Technical University of Graz, A-8010 Graz, Austria.

Author D. Haussler gratefully acknowledges the support of NSF grant IST-8317918. This work was conducted while author E. Welzl visited the University of Denver.

Abstract We present a new technique for half-space and simplex range query using $O(n)$ space and $O(n^\alpha)$ query time, where $\alpha < \frac{d(d-1)}{d(d-1)+1} + \gamma$ for all dimensions $d \geq 2$ and $\gamma > 0$. These bounds are better than those previously published for all $d \geq 2$. The technique uses random sampling to build a partition-tree structure. We introduce the concept of an ε-net for an abstract set of ranges to describe the desired result of this random sampling and give necessary and sufficient conditions that a random sample is an ε-net with high probability. We illustrate the application of these ideas to other range query problems.

1. Introduction

Rapid processing of geometric range queries has proven to be of fundamental importance in computational geometry, both as an end in itself and as a technique in the efficient solution of other geometric problems. Yao and Yao [*YAO 85*] have recently demonstrated that a wide variety of range query problems can be reduced to half-space query problems, the basic form of which may be described as follows: Given a set of n points in d-dimensional Euclidean space E^d, find a data structure that uses $O(n)$ storage such that the number of points in any given half-space can be determined quickly (i.e. in sublinear time). This problem is called the ***half-space enumeration problem***. A common variant of this problem is the ***reporting problem***, in which the set of points in the half-space is determined.

The first sublinear bounds for half-space enumeration queries were given by Willard [*WIL 82*], who showed that $O(n^\alpha)$ queries are possible in E^2 for $\alpha \cong 0.774$. Subsequently, Edelsbrunner and Welzl [*EDL 82*] improved this to $\alpha \cong 0.695$. In E^3, the first bound is Yao's [*YAO 83*] ($\alpha \cong 0.936$), which is followed by Dobkin and Edelsbrunner [*DOBa 84*] ($\alpha \cong 0.916$), Edelsbrunner and Huber [*EDL 84*] ($\alpha \cong 0.909$), Dobkin, Edelsbrunner and F. Yao [*DOBb 84*] ($\alpha \cong 0.899$). Shortly after Cole [*COL 85*] showed $\alpha \cong 0.977$ in E^4, Yao and Yao [*YAO 85*] gave a generalized version of this result, showing that $\alpha = \frac{\log(2^d - 1)}{d}$ can be achieved for all $d \geq 2$. This bound is the best published for $d \geq 4$. In this paper we exhibit a data structure that allows half-space enumeration queries in $O(n^\alpha)$ for $\alpha = \frac{d(d-1)}{d(d-1)+1} + \gamma$, for any $\gamma > 0$, which improves on previous bounds for all $d \geq 2$. Specific bounds are: $\alpha \cong 0.667$ (2 dimensions), $\alpha \cong 0.857$ (3 dimensions) and $\alpha \cong 0.923$ (4 dimensions). The technique also works for enumeration queries when ranges are simplices in d dimensions. Bounds for reporting are similar to those for enumeration, except that the number of points reported must be added to the time bound. It should be noted that better bounds are possible for reporting in two dimensions (that is, $O(\log n + t)$, where t is the number of points reported [*CHA 83*]), but these techniques only work for half-planes.

Our techniques are fundamentally similar to previous techniques employed for range queries. A partition tree is constructed so that a recursive divide-and-conquer strategy can be efficiently applied to any query. The main difference is that our construction is probabilistic, using random sampling to build each level of the partition tree. It is analogous to the technique used by Clarkson [*CLA 85*] to build efficient data structures for nearest neighbor queries.

The construction of the partition tree in E^2 can be described as follows. Given a set A of n data points in the plane, create a root node and choose at random a subset N of A of size v.[1] Form the line arrangement consisting of all lines defined by pairs of points in N. For each cell in this arrangement, create a child of the root that contains the number of points of A that lie in this cell. Now proceed recursively for each of these children, creating a subtree for the points in its cell until each cell contains less than v points.

This tree is queried in the usual manner. Given a half-plane determined by a line, the point counts from all cells at the first level of the tree that are completely contained in the half-plane are summed, and recursive calls are made for any cells that are cut by the line. The trick in establishing the time bound is to choose v such that with high probability

(1) the total number of points in all cells that are intersected by any line is less than εn for some small ε (where n is total number of points at the current level of recursion),

while keeping the total number of cells intersected reasonably small (this number is bounded by $O(v^2)$). Thus, in contrast to previous techniques, we do not rely on subdivisions of the points into parts of certain sizes. We show that it is enough to choose v to be $\left\lceil c\frac{1}{\varepsilon}(\log\frac{1}{\varepsilon}+\log\frac{1}{\delta})\right\rceil$, for some constant c, to get (1) with probability at least $1-\delta$ at any level of the tree. The key point is that for half-plane queries, the number v is essentially a constant that is independent of the size of the point set A, depending only on the parameters ε and δ. As ε approachs zero, we obtain the asymptotic results stated above with arbitrarily small γ, and as δ approaches zero, this happens with arbitrarily high probability. It should be noted that the partition tree is guaranteed to give the correct answer to any query independently of the choice of δ. It is only the time bound for queries that is achieved with probability depending on δ. On the other hand, this shows that for every point set and every $\gamma > 0$ there exists a partition tree that gives the claimed time bound. However, we have not been able to show that such a tree can be efficiently constructed (with probability 1).

In establishing the existence of the number v described above, we build on concepts due to Vapnik and Chervonenkis on uniformly approximating classes of events by their empirical distributions [*VAP 71*]. In extending their results, we introduce a new geometrical concept that may be of independent interest. Given a finite point set $A \subseteq E^d$ for some $d \geq 1$, a class of ranges R such that $r \subseteq E^d$ for all $r \in R$ and $\varepsilon \geq 0$, an *ε-net* of A for R is a set of points $N \subseteq A$ such that N contains a point in r for every $r \in R$ with $\frac{|A \cap r|}{|A|} > \varepsilon$. For example, if $\varepsilon = 0$ and R is the set of half-spaces, then the smallest ε-net of A for R is the set of all extreme points of A. It follows that when all points of A are extreme, the smallest 0-net of A for half-spaces is A itself. This cannot occur for $\varepsilon > 0$ and R the set of d-dimensional half-spaces. We show that for any $\varepsilon > 0$ and any finite point set $A \subseteq E^d$, there exists an ε-net of A for half-spaces with at most $\left\lceil \frac{8(d+1)}{\varepsilon}\log\frac{8(d+1)}{\varepsilon}\right\rceil$ points.

More generally, we characterize the classes of ranges for which there exists a function $f(\varepsilon)$ such that any finite point set A has an ε-net of size $f(\varepsilon)$, independently of the size of A. These are

[1] It actually suffices to make v independent random draws, even though this might not produce v distinct points. v is defined below.

precisely the classes of ranges with finite Vapnik-Chervonenkis dimension, known as Vapnik-Chervonenkis classes [*VAP 71*] [*DUD 78*], [*WEN 81*]. From this characterization result, it follows that if there exists any function $f(\varepsilon)$ such that any finite point set has an ε-net for R of size $f(\varepsilon)$, then in fact any finite point set has an ε-net for R of size at most $\left\lceil \frac{8d}{\varepsilon}\log\frac{8d}{\varepsilon} \right\rceil$, where d is the Vapnik-Chervonenkis dimension of R. Since the Vapnik-Chervonenkis dimension of the class of half-spaces in E^d is $d+1$, the above cited result is a special case of this theorem. Moreover, we show that if R has Vapnik-Chervonenkis dimension $d < \infty$ then for any $\varepsilon, \delta > 0$ and any finite point set A, if at least $\max\left(\frac{4}{\varepsilon}\log\frac{2}{\delta}, \frac{8d}{\varepsilon}\log\frac{8d}{\varepsilon}\right)$ points are drawn independently at random from A then these points form an ε-net of A for R with probability at least $1-\delta$. This latter result is used to obtain the number v above, used in the time bounds for our partition tree construction. Using the related notion of an *ε-approximation* (directly from [*VAP 71*]), we also exhibit trivial data structures of constant size that give approximate solutions to the enumeration problem for half-spaces in constant time.

Since the Vapnik-Chervonenkis dimension of half-spaces plus spherical ranges in d-dimensions is $d+1$ [*DUD 78*], the probabilistic aspects of Clarkson's RPO construction for nearest neighbor queries [*CLA 85*] can also be derived from the general properties of ε-nets and Vapnik-Chervonenkis classes. These concepts, in a slightly more general form, have also proven useful in the investigation of learning algorithms for concepts defined by geometrical regions in feature space [*BLU 85*]. It is our expectation that these concepts, along with the ε-approximation concept, will find other applications in computational geometry as well.

2. Geometric Fundamentals

For $d \geq 1$, let E^d denote d-dimensional Euclidean space. A point x in E^d is specified by its d coordinates in a Cartesian system.

A nonzero point x in E^d defines a hyperplane $h = \{y \in E^d : (x,y) = 1\}$, where (x,y) denotes the scalar product of x and y. The hyperplane h defines two open half-spaces in E^d, the positive half-space $h^+ = \{y \in E^d : (x,y) > 1\}$ and the negative half-space $h^- = \{y \in E^d : (x,y) < 1\}$. We will use h^* to denote one of these two open half-spaces determined by a hyperplane h, if we do not want to specify whether we mean the positive or the negative half-space. By $\bar{h}^+$, $\bar{h}^-$, and $\bar{h}^*$ we denote the closures of h^+, h^-, and h^*, respectively.

We use H_d, $H_d{}^+$, $H_d{}^-$, and $H_d{}^*$ to denote the set of all hyperplanes in E^d, the set of all positive half-spaces in E^d, the set of all negative half-spaces in E^d, and the set of all half-spaces in E^d, respectively. If N is a set of at least d points in general position in E^d, then $H_d(N)$ denotes the set of all hyperplanes that contain d of the points in N.

For $d \geq 0$, and $n \geq 0$ integers, $\Phi_d(n)$ is defined as follows: $\Phi_d(0) = 1$ for all $d \geq 0$, $\Phi_0(n) = 1$ for all $n \geq 0$, and $\Phi_d(n) = \Phi_d(n-1)+\Phi_{d-1}(n-1)$ for $d,n \geq 1$.

Proposition. $\Phi_d(n) = \sum_{k=0}^{d}\binom{n}{k}$ if $d < n$, otherwise $\Phi_d(n) = 2^n$. ▪

We assume familiarity with basic notions about arrangements of hyperplanes in E^d; in particular, we use the notion of a cell in an arrangement.

Proposition. Let H be an arrangement of n hyperplanes in E^d. Then the number of cells in this arrangement is at most $\Phi_d(n)$. The number of cells in this arrangement is equal to $\Phi_d(n)$ if the hyperplanes in H are in general position, i.e., no two hyperplanes in H are parallel and no $d+1$ hyperplanes in H have a common point. ▪

3. Range spaces, epsilon-nets and epsilon-approximations

In this section we introduce abstract range spaces and give upper bounds on the the number of points needed to form

ε-nets and ε-approximations for these spaces. The key concepts of this section are inspired by pioneering work of Vapnik and Chervonenkis [*VAP 71*].

Definition. A ***range space*** S is a pair (X,R), where X is a set and R is a set of subsets of X. Members of X are called ***elements*** or ***points*** of S and members of R are called ***ranges*** of S. S is ***finite*** if X is finite. ▪

The critical combinatorial parameter associated with a range space is its dimension, introduced in the following definition.

Definition. Let $S = (X,R)$ be a range space and let $A \subseteq X$ be a finite set of elements of S. Then $\Pi_R(A)$ denotes the set of all subsets of A that can be obtained by intersecting A with a range of S, i.e.

$\Pi_R(A) = \{A \cap r : r \in R\}$.

If $\Pi_R(A) = 2^A$, then we say that A is ***shattered*** by R. The ***Vapnik-Chervonenkis dimension*** of S (or simply the ***dimension*** of S) is the smallest integer d such that no $A \subseteq X$ of cardinality $d + 1$ is shattered by R. If no such d exists, we say the dimension of S is infinite. ▪

Examples. Consider the range space $S_1 = (E^1, H_1^+)$ consisting of the real line and all closed half-lines that are unbounded on the right. For any two points a and b with $a < b$, only the three subsets ϕ, $\{b\}$ and $\{a,b\}$ of $\{a,b\}$ can be formed by intersections with ranges in H_1^+, hence no 2 points of S_1 are shattered. Since it is obvious that a singleton set is shattered, this implies that the dimension of S_1 is 1. If we extend S_1 to $S'_1 = (E^1, H_1^*)$ by adding half-lines unbounded on the left, it is readily verified that 2 points can be shattered, but three points cannot, thus the dimension is 2.

These results generalize to higher dimensions, so that for any $d \geq 1$, (E^d, H_d^+) is of dimension d and (E^d, H_d^*) is of dimension $d + 1$. To see this, consider the dual image of a set A of n points in E^d. This yields a set of n hyperplanes that partition E^d into $\leq \Phi_d(n)$ cells, with equality whenever the points in A are in general position. Each of these cells corresponds to a unique intersection of a half-space in H_d^+ with A. Since $\Phi_d(n) = 2^n$ for $n = d$ and $\Phi_d(n) < 2^n$ for all $n > d$, this implies that the dimension of (E^d, H_d^+) is d. The other bound follows from an easy extension of this argument. Finally, it is clear that if, for example, $X = E^2$ and R is the set of all convex polygons, then there exist finite sets of points of every size that can be shattered (e.g. sets of points on the unit circle), hence the dimension of (X,R) is infinite. ▪

When (X,R) is of finite dimension, Dudley calls R a ***Vapnik-Chervonenkis Class*** (VCC) [*DUD 78*] [*WEN 81*]. Dudley's notion of the ***Vapnik-Chervonenkis number*** of R corresponds to the dimension of (X,R) plus one. Translating into our terminology, Dudley shows that whenever (X,R) is of finite dimension, then (X,R_k) is also of finite dimension, where R_k is the set of all Boolean combinations formed from at most k ranges in R. Thus for example, since the set C_k of k-gons in E^d for fixed $k > d$ is formed by k-fold intersections of half-spaces, (E^d, C_k) is of finite dimension for any finite k. We give bounds on the dimension of this and related spaces below. Dudley and Wenocur also prove more general results that imply e.g. that the range space formed by the set of all half-spaces bounded by polynomial curves of fixed degree also has finite dimension.

The function $\Phi_d(n)$ plays a fundamental role in all range spaces of finite dimension. The following lemma and theorem constitute a slightly stronger version of the result given in Lemma 1 of [*VAP 71*].

Lemma 3.1. Let (X,R) be a finite range space of dimension d with $|X| = n$. Then $|R| \leq \Phi_d(n)$.

Proof. The assertion is trivially true for $d = 0$ and $n = 0$. Assume the assertion is true for any finite range space of dimension at most $d - 1$, and for any range space of dimension d with at most $n - 1$ elements, for some $d \geq 1$ and $n \geq 1$.

Let (X,R) be a range space of dimension d with $|X| = n$ and $x \in X$. Consider the range spaces $S-x = (X - \{x\}, R-x)$, where $R-x = \{r - \{x\} : r \in R\}$ and $S^{(x)} = (X - \{x\}, R^{(x)})$, where $R^{(x)} = \{r \in R : x \notin r, \ r \cup \{x\} \in R\}$. Obviously $S-x$ is of dimension at most d; hence, by assumption, $|R-x| \le \Phi_d(n-1)$. We show that $S^{(x)}$ is of dimension at most $d-1$.

Let A be a subset of $X - \{x\}$ that can be shattered by $R^{(x)}$. Then it is easy to see that $A \cup \{x\}$ can be shattered by R. (For $A' \subseteq A$ there is an $r \in R^{(x)}$ with $A' = A \cap r$. Since $x \notin r$, $A' = (A \cup \{x\}) \cap r$ and $A' \cup \{x\} = (A \cup \{x\}) \cap (r \cup \{x\})$, where both r and $r \cup \{x\}$ are in R.) Since $A \cup \{x\}$ can be shattered by R, $|A \cup \{x\}| \le d$, so $|A| \le d-1$. Thus $S^{(x)}$ is of dimension at most $d-1$.

Since $S^{(x)}$ is of dimension at most $d-1$, by assumption, $|R^{(x)}| \le \Phi_{d-1}(n-1)$. Observing that $|R| = |R-x| + |R^{(x)}|$, this yields $|R| \le \Phi_d(n-1) + \Phi_{d-1}(n-1) = \Phi_d(n)$. ▪

Taking X to be a set of n points in general position in E^d and R the set of all intersections of X with positive half-spaces shows that the bound in the above lemma is the best possible. This lemma extends to arbitrary range spaces of finite dimension as follows.

Theorem 3.2. Let (X,R) be a range space of dimension d. For every finite subset A of X, $|\Pi_R(A)| \le \Phi_d(|A|)$.

Proof. We need only observe that if $A \subseteq X$, where (X,R) is a range space of dimension d, then $(A, \Pi_R(A))$ is a finite range space of dimension at most d, and use the above lemma. ▪

The original motivation for the study of Vapnik-Chervonenkis classes was to determine the classes of sets whose probability measures could be uniformly approximated by random sampling. Since we are only concerned with finite point sets in this paper, we make the following

Definition. Let (X,R) be a range space and A a finite subset of X. For any $\varepsilon \ge 0$ and $V \subseteq A$, V is an *ε-approximation of* A (for R) if for all $r \in R$ $\left| \frac{|A \cap r|}{|A|} - \frac{|V \cap r|}{|V|} \right| \le \varepsilon$. ▪

Theorem 3.3. ([VAP 71]) Let (X,R) be a range space of dimension d, $A \subseteq X$ be a finite set and $\varepsilon, \delta > 0$. Then any random sample V of A formed by at least m independent draws from A is an ε-approximation of A for R with probability at least $1-\delta$ for any

$$m \ge \frac{16}{\varepsilon^2}\left(d \ln \frac{16d}{\varepsilon^2} + \ln \frac{4}{\delta}\right). \ ▪$$

Corollary 3.4. For every $d \ge 1$, if (X,R) is a range space of dimension d, A is a finite subset of X and $\varepsilon > 0$, then there exists an ε-approximation V of A with respect to R, with

$$|V| \le \left\lceil \frac{16}{\varepsilon^2}\left(d \ln \frac{16d}{\varepsilon^2} + 2\right) \right\rceil ▪$$

Example. Let A be a set of points in E^2 of cardinality n. Since the dimension of positive half-planes is 2, the above formula implies that there exists a subset V of A of size at most 4,376,345 such that for every positive half-plane h, $|A \cap h|$ and $\frac{|V \cap h|}{|V|} n$ differ by at most 1% of n. ▪

This leads to a simple data structure for point sets that answers half-plane enumeration queries to any desired accuracy (in proportion to the size n of the point set) in constant time. We simply find the set V as above, count the fraction of points of V in the half-plane and multiply by n. Moreover, the set V can be found with arbitrarily high probability by simply drawing a random sample of the point set. Since Theorem 3.3 holds for any range space of finite dimension, similar results hold when half-planes are replaced by circles, triangles, etc, as well as their higher dimensional counterparts. However, the constants involved can be quite large.

Our main goal is to find algorithms for half-space queries that give exact answers using linear space and sublinear time. To this end, we introduce a new concept related to that of an ε-approximation.

Definition. Let (X,R) be a range space, A a finite subset of X and $\varepsilon \ge 0$. Then $R_{A,\varepsilon}$ denotes the set of all $r \in R$ that contain a fraction of the points in A of size greater than ε, i.e. such that $\frac{|A \cap r|}{|A|} > \varepsilon$. A subset N of A is an ***ε-net*** *of* A (for R) if N contains a point in each $r \in R_{A,\varepsilon}$. ▪

Example. Consider n points A on a circle in E^2. For any $\varepsilon > 0$, an ε-net of A for half-planes can be found by choosing a subset N of A such that among any $\lfloor \varepsilon n \rfloor + 1$ consecutive points on the circle at least one point is in N. This can clearly be done using at most $\left\lceil \frac{1}{\varepsilon} \right\rceil$ points. ▪

We now give bounds on the sizes of ε-nets for arbitrary finite dimensional range spaces analogous to those given above for ε-approximations. The following will make our arguments clearer.

Notation. For any finite set A and $m \geq 1$, A^m denotes the m-fold cross product of A. A vector $x_1 \cdots x_m \in A^m$ will be denoted $\bar{x}$ when m is clear from the context. Similarly, $\bar{y}$ denotes $y_1 \cdots y_m$. For any $S \subseteq A^m$, $P_A{}^m(S)$ denotes the probability that a vector in S is obtained in m independent draws with replacement from A, i.e. $P_A{}^m(S) = \frac{|S|}{|A|^m}$. ▪

For the following two lemmas, let (X,R) be a fixed range space and A be a fixed finite subset of X.

Definition. For any $m \geq 1$ and $\varepsilon > 0$, let $Q_\varepsilon{}^m$ be the set of all m-vectors whose elements do not form an ε-net of A, i.e. $Q_\varepsilon{}^m = \{\bar{x} \in A^m :$ there exists $r \in R_{A,\varepsilon}$ such that $x_i \notin r$, $1 \leq i \leq m\}$. Let $J_\varepsilon{}^{2m} = \{\bar{x}\bar{y} \in A^{2m}$ (where $\bar{x},\bar{y} \in A^m$) : there exists $r \in R_{A,\varepsilon}$ such that $x_i \notin r$, $1 \leq i \leq m$ but $y_i \in r$ for at least $\frac{\varepsilon m}{2}$ indices i, $1 \leq i \leq m\}$. ▪

Lemma 3.5. $P_A{}^m(Q_\varepsilon{}^m) < 2P_A{}^{2m}(J_\varepsilon{}^{2m})$ for all $\varepsilon > 0$ and $m \geq \frac{8}{\varepsilon}$.

Proof. For $r \in R_{A,\varepsilon}$ let $S_r = \{\bar{y} \in A^m : y_i \in r$ for at least $\frac{\varepsilon m}{2}$ indices i, $1 \leq i \leq m\}$. We first claim that $P_A{}^m(S_r) > \frac{1}{2}$ for all $r \in R_{A,\varepsilon}$. To establish this, we show that $P_A{}^m(\bar{S}_r) < \frac{1}{2}$, where $\bar{S}_r = A^m - S_r$. Since $P_A(r) \geq \varepsilon$ for each $r \in R_{A,\varepsilon}$ and $\bar{y} \in \bar{S}_r$ only if $y_i \in r$ for fewer than $\frac{\varepsilon m}{2}$ indices i, $P_A{}^m(\bar{S}_r)$ is maximized as $P_A(r)$ approaches ε. In this case, for random $\bar{y} \in A^m$ the expected number of indices i such that $y_i \in r$ is εm and the variance is $\varepsilon(1-\varepsilon)m$. Thus for each $\bar{y} \in \bar{S}_r$, the number of y_i's in r differs by at least $\frac{\varepsilon m}{2}$ from the expected value. Hence by Chebyshev's inequality $P_A{}^m(\bar{S}_r) \leq \frac{\varepsilon(1-\varepsilon)m}{(\frac{\varepsilon m}{2})^2} < \frac{4}{\varepsilon m} \leq \frac{1}{2}$, since $m \geq \frac{8}{\varepsilon}$, establishing the claim.

Now consider a fixed $\bar{x} \in Q_\varepsilon{}^m$. By definition, there exists $r \in R_{A,\varepsilon}$ such that $x_i \notin r$, $1 \leq i \leq m$. From the above, it follows that for more than half of the $\bar{y} \in A^m$, $\bar{y} \in S_r$, hence $\bar{x}\bar{y} \in J_\varepsilon{}^{2m}$. Hence $P_A{}^{2m}(J_\varepsilon{}^{2m}) > \frac{1}{2}P_A{}^m(Q_\varepsilon{}^m)$. ▪

Lemma 3.6. If (X,R) is of finite dimension d, $P_A{}^{2m}(J_\varepsilon{}^{2m}) < \Phi_d(2m)2^{\frac{-\varepsilon m}{2}}$.

Proof. For each j, $1 \leq j \leq (2m)!$, let π_j be a distinct permutation of the indices $1, \ldots, 2m$. For each $\bar{x} \in A^{2m}$, let $\Theta(\bar{x}) = |\{j : \pi_j(\bar{x}) \in J_\varepsilon{}^{2m}\}|$. It is easily verified that $P_A{}^{2m}(J_\varepsilon{}^{2m}) \leq \max_{\bar{x} \in A^{2m}} \left\{\frac{\Theta(\bar{x})}{(2m)!}\right\}$. Consider a fixed $\bar{x} \in A^{2m}$. Let S be the set of distinct elements of A that appear in $\bar{x}$. Since $|S| \leq 2m$ and (X,R) is of dimension d, by Theorem 3.2 there are at most $\Phi_d(2m)$ distinct subsets of S induced by intersections with $r \in R_{A,\varepsilon}$. Each such subset T of S is a witness to the fact that certain permutations of $\bar{x}$ are in $J_\varepsilon{}^{2m}$. Specifically, whenever all occurrences of members of T appear in the second half of $\pi_j(\bar{x})$, and there are at least $\frac{\varepsilon m}{2}$ such occurrences, then $\pi_j(\bar{x}) \in J_\varepsilon{}^{2m}$, otherwise $\pi_j(\bar{x}) \notin J_\varepsilon{}^{2m}$. However, for a given T this can occur in only a small fraction of all permutations of $\bar{x}$. In particular, if there are l occurrences of members of T in $\bar{x}$, then T is a witness for at most

$$\frac{\binom{m}{l}}{\binom{2m}{l}} = \frac{m(m-1)\ldots(m-l+1)}{2m(2m-1)\ldots(2m-l+1)} \leq 2^{-l} \leq 2^{\frac{-\varepsilon m}{2}}$$

of all permutations of $\bar{x}$. It follows that $\frac{\Theta(\bar{x})}{(2m)!} \leq \Phi_d(2m)2^{\frac{-\varepsilon m}{2}}$. ▪

Directly from the above two lemmas we get the following

Theorem 3.7. If (X,R) has finite dimension d and $m \geq \frac{8}{\varepsilon}$ then

$$P_A{}^m(Q_\varepsilon{}^m) \le 2\Phi_d(2m)2^{\frac{-\varepsilon m}{2}}$$

■

Corollary 3.8. For any (X,R) of dimension $d < \infty$, finite $A \subseteq X$ and $\varepsilon,\delta > 0$, if N is the set of distinct elements of A obtained by $m \ge \max\left\{\frac{4}{\varepsilon}\log\frac{2}{\delta}, \frac{8d}{\varepsilon}\log\frac{8d}{\varepsilon}\right\}$ random independent draws from A, then N is an ε-net of A for R with probability at least $1-\delta$.

Proof. It is easily shown that for m of this size, the bound given in the Theorem 3.7 is less than δ (see [*BLU 85*], Lemma 7). Hence the sample N will be an ε-net with probability at least $1-\delta$. ■

Corollary 3.9. For any range space (X,R) the following are equivalent.
(i) (X,R) has finite dimension.
(ii) There exists a function $f(\varepsilon)$ such that for all $\varepsilon > 0$ and all finite $A \subseteq X$, there exists an ε-net for A (for R) with at most $f(\varepsilon)$ elements, independently of the size of A.
(iii) There exists d such that for all $\varepsilon > 0$ and all finite $A \subseteq X$, there is an ε-net for A (for R) with at most $\left\lceil\frac{8d}{\varepsilon}\log\frac{8d}{\varepsilon}\right\rceil$ elements.

Proof. From the above corollary, (i) implies (iii), and it is obvious that (iii) implies (ii). To see that (ii) implies (i), suppose to the contrary that (ii) holds, but (X,R) has infinite dimension. Then for each $n \ge 1$ we can find a subset A_n of X of size n that is shattered by R. To obtain an ε-net of A_n it is clear that at least $\lfloor(1-\varepsilon)n\rfloor$ points must be used, otherwise there is a subset of A_n in $R_{A,\varepsilon}$ with no points in it. Since this function grows with n, this contradicts (ii). ■

Example. For any set A of points in E^2 there is a subset N of A of size at most 17,031 such that every positive half-plane that contains at least 1% of the points of A contains at least one point in N. Since the Vapnik-Chervonenkis dimension of the set of all triangles in E^2 is 7, if the size of N is increased to 69,727, the same result holds for triangles. Better estimates can be obtained by using Theorem 3.7 directly, setting the bound given there to be less than 1. Even here no attempt has been made to minimize the constants, so it is likely that these results actually hold for considerably smaller numbers. ■

4. Range search data structures

We now describe the data structure we use for the half-space range search enumeration problem for point sets in E^d. This data structure can be easily extended to the reporting problem and to simplices instead of half-spaces. In order to simplify the exposition, we restrict ourselves to point sets in general position in E^d.

Definition. Let $d \ge 2$ and $v \ge d$ be integers and $0 \le \varepsilon \le 1$. A rooted tree T is called an *(ε,v)-partition tree* for a finite point set A in general position in E^d, if the following hold.
(1) Every node p of T corresponds to some open region $reg(p)$ of E^d.
(2) The root r of T corresponds to $reg(r) = E^d$.
(3) Let p be a node of T that corresponds to $reg(p)$.
(3.1) If $|A \cap reg(p)| < v$, then p is a *leaf of* T that has associated:
  (i) The set $set(p) = A \cap reg(p)$.
  (ii) The number $a(p) = |A \cap reg(p)|$.
(3.2) If $|A \cap reg(p)| \ge v$, then p is an *internal node* of T that has associated:
  (i) A set $set(p) \subseteq A \cap reg(p)$ of cardinality v.
  (ii) The arrangement $arr(p)$ of all hyperplanes that contain d of the points in $set(p)$.
  (iii) For each cell f in $arr(p)$ with $f \cap reg(p) \cap A \ne \phi$, a child p_f of p with $reg(p_f) = f \cap reg(p)$.
  (iv) The number $a(p) = |A \cap reg(p)|$.
(4) For every internal node p in T and every hyperplane h in E^d

$$\sum_{f \in F} |f \cap reg(p) \cap A| \le \varepsilon |reg(p) \cap A|,$$

where F is the set of all cells f in $arr(p)$ with $f \cap h \ne \phi$. ■

Condition (4) is trivially satisfied for $\varepsilon = 1$. For smaller ε, we will see below that this condition leads to sublinear enumeration search times for half-space queries.

The objects associated with the nodes the tree are used to support half-space queries. The search for a half-space h^* will be performed by a procedure $ENUM(p,h^*)$, p a node in the tree. We use a global variable $ANSW$ that is set to zero in the beginning and that will hold the integer $|A \cap h^*|$ after the return from the call $ENUM(r,h^*)$, where r is the root of the (ε,v)-partition tree storing the point set A. The procedure $ENUM(p,h^*)$ can be described as follows:

$ENUM(p,h^*)$:
  $ANSW := ANSW + |set(p) \cap h^*|$;
    if p is an internal node *do:*
      for each cell f in $arr(p)$ for which child p_f exists *do:*
        if $f \subseteq h^*$, *then* $ANSW := ANSW + a(p_f)$;
        else if f is intersected by h, *then* $ENUM(p_f,h^*)$.

Note that we make the call $ENUM(p_f,h^*)$ under the condition "f is intersected by h" and not under the condition "$reg(p_f)$ is intersected by h", thus avoiding the need to store $reg(p)$ for each node p in T.

We make the following observations.

Observation 4.1. Let $QT(p,h^*)$ be the query time required by the call $ENUM(p,h^*)$ in an (ε,v)-partition tree T storing a point set in E^d.

(1) If p is an internal node then

$$QT(p,h^*) \le O(\Phi_d(\tbinom{v}{d})) + \sum_{f \in F} QT(p_f,h^*)),$$

where F is the set of all cells f in $arr(p)$ with $f \cap h \ne \phi$ for which p_f exists.

(2) If p is a leaf then $QT(p,h^*) \le O(v)$. ▪

Observation 4.2. An (ε,v)-partition tree stores a set of n points in general position in E^d in $O(v^{d^2}n)$ space. ▪

Lemma 4.3. An (ε,v)-partition tree T, $0 < \varepsilon < 1$, storing a set A of n points in E^d in general position, supports the computation of $|A \cap h^*|$ for every half-space h^* in $O(v^{d^8}n^{\alpha})$ time with

$$\alpha = \frac{\log_{\frac{1}{\varepsilon}} \Phi_{d-1}(\tbinom{v}{d})}{\log_{\frac{1}{\varepsilon}} \Phi_{d-1}(\tbinom{v}{d})+1}.$$

Proof. By Observation 4.1, for a node p in T and a half-space h^*, the time required by the call $ENUM(p,h^*)$ is

$$QT(p,h^*) \le O(v^{d^8}) + \sum_{f \in F} QT(p_f,h^*),$$

where F is the set of all cells f in $arr(p)$ with $f \cap h \ne \phi$ for which p_f exists. Moreover we know that

$$|F| \le \Phi_{d-1}(\tbinom{v}{d}) \text{ and}$$
$$\sum_{f \in F} |reg(p_f) \cap A| \le \varepsilon |reg(p) \cap A|.$$

For a nonnegative integer n, let $Q(n)$ be the maximal time required by a call $ENUM(p,h^*)$ for any p in any (ε,v)-partition tree storing a set A in E^d and any half-space h^*, where $|reg(p) \cap A| \le n$. Then

$$Q(n) \le \max(O(v^{d^8}) + \sum_{i=1}^{m} Q(n_i))$$

where $m = \Phi_{d-1}(\tbinom{v}{d})$ and the maximum is taken over all $(n_1,n_2,\ldots,n_m)$ with $\sum_{i=1}^{m} n_i \le \varepsilon n$. Since Q is a convex function,

$$\sum_{i=1}^{m} Q(n_i) \le mQ\left(\left\lceil \frac{\sum_{i=1}^{m} n_i}{m} \right\rceil\right),$$

which implies the claimed time bound. ▪

It is clear that (ε,v)-partition trees do not exist for every point set A in E^d and every ε and v. For example, a $(0,v)$-partition tree exists for a set A in E^d (in general position) if and only if $v \ge |A|$. In the next few lemmas we show that for $\varepsilon > 0$, (ε,v)-partition trees always exist with v "reasonably small" compared to $\frac{1}{\varepsilon}$.

Lemma 4.4. Let N be a set of $v \ge d+1$ points in E^d in general position and let h_0 be a hyperplane in E^d.

(1) If $N \subseteq \bar{h}_0^*$ for $h_0^* = h_0^+$ or $h_0^* = h_0^-$, then there exist k hyperplanes in $H_d(N)$, $h_1, h_2, \cdots, h_k$, where $k \le d$, such that for each i, $1 \le i \le k$, there exists $h_i^* \in \{h_i^+, h_i^-\}$ with $N \subseteq \bar{h}_i^*$, and

$$h_0 \subseteq E^d - \bigcap_{i=1}^{k} h_i^*.$$

(2) If both $N_1 = h_0^+ \cap N$ and $N_0 = h_0^- \cap N$ are nonempty, then there exist k hyperplanes in $H_d(N)$, $h_1, h_2, \cdots, h_k$, where $k \le d+1$, such that for each i, $1 \le i \le k$, there exists $h_i^1 \in \{h_i^+, h_i^-\}$ with $h_i^1 \cap N \subseteq N_1$ and $h_i^0 \in \{h_i^+, h_i^-\}$ with $h_i^0 \cap N \subseteq N_0$ and

$$h_0 \subseteq \bigcup_{i=1}^{k} \bar{h}_i^1 - \bigcap_{i=1}^{k} h_i^1.$$

Proof. We outline here only the idea for (2).

Let H be the set of hyperplanes h in $H_d(N)$ such that $h^1 \cap N \subseteq N$, for one half-space h^1 in $\{h^+, h^-\}$ and $h^0 \cap N = N_0$ for the other half-space h^0 in $\{h^+, h^-\}$. Then it is easy to see that

$$h \subseteq \bigcup_{h \in H} \bar{h}^1 - \bigcap_{h \in H} h^1.$$

Now the assertion can be seen as the dual foundation of Caratheodry's Theorem (see [*GRU 67*], Theorem 2.3.5), which states that if a point x is in the convex hull of a set A in E^d, then there exists a subset A' of A such that $|A'| \le d+1$ and x is in the convex hull of A'. ▪

Part (2) of the above lemma leads to the following definition.

Definition. Let H^* be a set of open half-spaces in E^d. Then the *corridor* defined by H^*, denoted $corr(H^*)$, is the open region of the form

$$corr(H^*) = \bigcup_{h^* \in H^*} h^* - \bigcap_{h^* \in H^*} \bar{h}^*$$

If H^* is finite, and $k \ge |H^*|$ then we call $corr(H^*)$ a *k-corridor*. ▪

We obtain upper bounds on the Vapnik-Chervonenkis dimension of k-corridors in E^d from the following general result.

Lemma 4.5. Assume $k \ge 1$ and (X,R) is a range space of dimension $d \ge 2$. Let R' be the set of all sets of the form $\bigcup_{i=1}^{k} r_i - \bigcap_{i=1}^{k} r_i$, where r_i is a range in R, $1 \le i \le k$. Then (X,R') has dimension less than $2dk\log(dk)$.

Proof. Clearly we may assume that $k \ge 2$. Consider a finite set $A \subseteq X$ with $|A| = m \ge 3$. By Theorem 3.2, $|\Pi_R(A)| \le \Phi_d(m)$. Every set in $\Pi_{R'}(A)$ is of the form $\bigcup_{i=1}^{k} r_i - \bigcap_{i=1}^{k} r_i$, with $r_i \in \Pi_R(A)$, $1 \le i \le k$. This shows that
$|\Pi_{R'}(A)| \le |\Pi_R(A)|^k \le (\Phi_d(m))^k < m^{dk}$.
Hence, if $m^{dk} \le 2^m$, then A cannot be shattered by R' and the dimension of (X,R') is at most $m-1$. It is easy to show that $m^{dk} \le 2^m$ for $dk \ge 2$ and $m = \lceil 2dk\log(dk) \rceil$. ▪

Corollary 4.6. The Vapnik-Chervonenkis dimension of the set of all k-corridors in E^d is at most $2(d+1)k\log((d+1)k)$. ▪

Lemma 4.7. Let A be a set of n points in general position in E^d, let N be an ε-net of A for $(d+1)$-corridors and let h_0 be a hyperplane. If F is the set of all cells in the arrangement formed by $H_d(N)$ that are intersected by h, then

$$\text{(i) } |F| \le \Phi_{d-1}\left(\binom{|N|}{d}\right) \text{ and}$$
$$\text{(ii) } \sum_{f \in F} |f \cap A| \le \varepsilon n.$$

Proof. Part (i) is easily verified. For Part (ii), assume first that there are points in N on either side of h_0 and let $h_1, h_2, \dots, h_k$, where $k \le d+1$, be hyperplanes in $H_d(N)$ as they are described in Lemma 4.5 (2), i.e., $h_i^1 \cap N \subseteq h_0^+ \cap N$ for one half-space h_i^1 determined by h_i and $h_i^0 \cap N \subseteq h_0^- \cap N$ for the other half-space determined by h_i. Then h_0 is contained in the closure of $C = \bigcup_{i=1}^{k} h_i^1 - \bigcap_{i=1}^{k} \bar{h}_i^1$, and C contains all cells in the arrangement of $H_d(N)$ that are intersected by h_0. Since $C \cap N = \phi$ and N is an ε-net of A for $(d+1)$-corridors, $|C \cap A| \le \varepsilon n$. When either $h_0^+ \cap N$ or $h_0^- \cap N$ is empty, the result is established by a similar reasoning using Lemma 4.5 (1). ▪

Lemma 4.8. For each $d \ge 2$ there is a constant $c_d > 0$ such that an (ε, v)-

partition tree exists for every ε and v, where $0 < \varepsilon < 1$ and $v \geq \left\lceil c_d \frac{1}{\varepsilon} \log \frac{1}{\varepsilon} \right\rceil$, for every finite point set A in E^d in general position.

Proof. Let T be a $(1,v)$-partition tree for a finite point set A in E^d and let $0 < \varepsilon < 1$. From Lemma 4.7, we know that if for all internal nodes p in T, $set(p)$ is an ε-net of $(reg(p) \cap A)$ for $(d+1)$-corridors, then T is an (ε,v)-partition tree for A. By Corollaries 3.9 and 4.6, there exists a constant c_d for each $d \geq 0$, such that for all ε, $0 < \varepsilon < 1$, there exists an ε-net for $(d+1)$-corridors of size $\leq \left\lceil c_d \frac{1}{\varepsilon} \log \frac{1}{\varepsilon} \right\rceil$. ▪

Theorem 4.9. Let $\gamma > 0$ and $d \geq 2$ be fixed. For every set A of n points in E^d there exists an $O(n)$ space data structure that supports the computation of $|A \cap h^*|$ for every half-space h^* in $O(n^\alpha)$ time with $\alpha = \frac{d(d-1)}{d(d-1)+1} + \gamma$. (The constants for space and time depend on γ and d.)

Proof. Consider an (ε,v)-partition tree T for a finite point set A in E^d with $v = \left\lceil c_d \frac{1}{\varepsilon} \log \frac{1}{\varepsilon} \right\rceil$ where c_d is a constant such that such an (ε,v)-partition tree exists for all $\varepsilon > 0$. Linear space complexity follows from Observation 4.1. By Lemma 4.3, the time complexity for a query in T is $O(v^{d^2} n^{\alpha_\varepsilon})$ where

$$\alpha_\varepsilon = \frac{\log_{\frac{1}{\varepsilon}} \Phi_{d-1}\left(\binom{v}{d}\right)}{\log_{\frac{1}{\varepsilon}} \Phi_{d-1}\left(\binom{v}{d}\right) + 1} \leq \frac{\log_{\frac{1}{\varepsilon}} v^{d(d-1)}}{\log_{\frac{1}{\varepsilon}} v^{d(d-1)} + 1} =$$

$$\frac{d(d-1)\log_{\frac{1}{\varepsilon}} c_d \frac{1}{\varepsilon}\log\frac{1}{\varepsilon}}{d(d-1)\log_{\frac{1}{\varepsilon}} c_d \frac{1}{\varepsilon}\log\frac{1}{\varepsilon} + 1} =$$

$$\frac{d(d-1)(1 + \log_{\frac{1}{\varepsilon}} c_d + \log_{\frac{1}{\varepsilon}} \log\frac{1}{\varepsilon})}{d(d-1)(1 + \log_{\frac{1}{\varepsilon}} c_d + \log_{\frac{1}{\varepsilon}} \log\frac{1}{\varepsilon}) + 1}.$$

Now we observe that $\log_{\frac{1}{\varepsilon}} c_d \to 0$ and $\log_{\frac{1}{\varepsilon}} \log \frac{1}{\varepsilon} \to 0$ as $\varepsilon \to 0$. Thus for any $\gamma > 0$,

$$\alpha_\varepsilon \leq \frac{d(d-1)}{d(d-1)+1} + \gamma$$

for small enough ε. ▪

The extension of the result for half-space range search to simplex range search can be seen as follows: a simplex s^* in E^d is defined by $d+1$ hyperplanes. Thus the boundary s of s^* is the subset of the union of $d+1$ hyperplanes. If we use now an (ε,v)-partition tree to answer a query for simplex s^* in E^d in the obvious way, then we recur in each cell intersected by s, and these cells contain at most a total of $(d+1)\varepsilon$ times the current number of points. Just as in the case of half-spaces, the constant $d+1$ is absorbed as ε approaches 0.

Finally, we mention the implication of the above result to a purely geometric problem as it was raised by Edelsbrunner [*EDL 85*]. The known bounds to this problem are the same as for the half-space range enumeration problem, as these problems are closely related. Thus we have also improved the exponent in this result for all dimensions d, $d \geq 2$.

Theorem 4.10. For each $d \geq 1$ and $\gamma > 0$ there exist constants k_1 and k_2 such that for every finite point set A and E^d there is a cell complex $C(A)$ that partitions E^d and has the following properties:

(i) each cell of $C(A)$ is a d-dimensional convex polytope.

(ii) no cell contains a point of A in its interior.

(iii) the number of cells in $C(A)$ is at most $k_1|A|$.

(iv) the maximum number of cells in $C(A)$ intersected by an arbitrary hyperplane is at most $k_2|A|^\alpha$, where $\alpha = \frac{d(d-1)}{d(d-1)+1} + \gamma$.

Proof. Consider an (ε,v)-partition tree T for A in E^d that realizes $O(n^\alpha)$ query time. Then the regions $reg(p)$, p a leaf in T, form a linear size cell complex that obviously satisfies (i), (iii) and (iv) for appropriate constants. However, since every region considered contains at most v points, it is also easy to ensure property (ii). ▪

Open Problems

Since the above results on Vapnik-

Chervonenkis dimension and ε-nets are quite general, we suspect they apply to many types of range queries beyond those we have explored here.

Note also that while the existence of "small" ε-nets for spaces of finite dimension plays a crucial role in our results, we give only probabilistic algorithms for constructing them. Efficient deterministic algorithms for constructing such nets remain to be determined. Here there might be some trade-offs in size versus time of computation. Perhaps more significantly, we have yet to determine the size of the smallest nets possible for many natural range spaces of finite dimension.[2]

Acknowledgement

We would like to thank Herbert Edelsbrunner for valuable discussions of these ideas and for his help in the preparation of this manuscript and Judson McCraine for pointing out errors in the numerical estimates given in an earlier draft.

References

[BLU 85] A. Blumer, A. Ehrenfeucht, D. Haussler, M. Warmuth, "Classifying learnable geometric concepts with the Vapnik-Chervonenkis dimension," *Proc. 18th Symp. on Theory of Computation*, (1986) to appear.

[CHA 83] B. Chazelle, L. Guibas and D.T. Lee, "The power of geometric duality," *Proc. 24th Symp. on Foundations of Comp Sci.*, 1983, 217-225.

[CLA 85] K. Clarkson, "A probabilistic algorithm for the post office problem," *Proc. 17th Symp. on Theory of Computation*, (1985) 175-185.

[COL 85] R. Cole, "Partitioning point sets in 4 dimensions," *Proc. Colloq. on Autom., Lang. and Progr., (1985)*, 111-119, Lecture Notes in Comput. Sci. 194, Springer, Berlin.

[DOBa 84] D. Dobkin and H. Edelsbrunner, "Organizing points in two and three dimensions," *Tech. Rep. F130, Technische Universitat Graz, 1984.*

[DOBb 84] D. Dobkin, H. Edelsbrunner and F. Yao, "A 3-space partition and its applications", manuscript.

[DUD 78] R.M. Dudley, "Central limit theorems for empirical measures," *Ann. Prob.*, 6 (6) (1978) 899-929.

[EDL 82] H. Edelsbrunner and E. Welzl, "Halfplane range search in linear space and $O(n^{0.695})$ query time," *Inf. Proc. Let.*, to appear.

[EDL 84] H. Edelsbrunner and F. Huber, "Dissecting sets of points in two and three dimensions," Tech. Rep. F138, Technische Universitat Graz, 1984. [EDL 85] H. Edelsbrunner, Open problem in *Bull. EATCS* 26 (1985).

[GRU 67] Branko Gruenbaum, *Convex Polytopes*, Interscience (1967)

[VAP 71] V.N. Vapnik and A.YA. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Th. Prob. and its Appl.* 16 (2) (1971) 264-280.

[WEN 81] R.S. Wenocur and R.M. Dudley, "Some special Vapnik-Chervonenkis classes," *Discrete Math.*, 31 (1891) 313-318.

[WIL 82] D. Willard, "Polygon retrieval," *SIAM J. Comput.*, 11 (1982) 149-165.

[YAO 83] F. Yao, "A 3-space partition and its applications," *Proc. 15th Symp. on Theory of Computation*, (1983) 258-263.

[YAO 85] A. Yao and F. Yao, "A general approach to d-dimensional geometric queries," *Proc. 17th Symp. on Theory of Computation*, (1985) 163-169.

[2] For half-planes in E^2 we can efficiently construct ε-nets of size $O(\frac{1}{\varepsilon})$ for any finite point set. However, our construction does not generalize to higher dimensions.