A non-linear lower bound for planar epsilon-nets

Noga Alon *

Abstract

We show that the minimum possible size of an ϵ-net for point objects and line (or rectangle)-ranges in the plane is (slightly) bigger than linear in $1/\epsilon$. This settles a problem raised by Matoušek, Seidel and Welzl in 1990.

1 Introduction

A *range space* S is a pair (X, R), where X is a (finite or infinite) set and R is a (finite or infinite) family of subsets of X. The members of X are called *points* and those of R are called *ranges*. If A is a subset of X then $P_R(A) = \{r \cap A : r \in R\}$ is the *projection* of R on A. In case this projection contains all subsets of A we say that A is *shattered*. The *Vapnik-Chervonenkis* dimension (or VC-dimension) of S, denoted by $VC(S)$, is the maximum cardinality of a shattered subset of X. If there are arbitrarily large shattered subsets then $VC(S) = \infty$.

For a finite set of points A in a range space, a subset $N \subset A$ is an *ϵ-net* for A if any range $r \in R$ satisfying $|r \cap A| \geq \epsilon|A|$ contains at least one point of N.

A well known result of Haussler and Welzl [10], following earlier work of Vapnik and Chervonenkis [20], asserts that for any n and $\epsilon > 0$, any set of size n in a range space of VC-dimension d contains an ϵ-net of size at most $O(\frac{d}{\epsilon}\log(1/\epsilon))$. See also [11] for a proof of a $(1 + o(1))(\frac{d}{\epsilon}\log(1/\epsilon))$ upper bound, and [13], [15], [3] for more details.

As shown in [11] there are known cases in which for fixed d the size of the smallest possible ϵ-net for a given set cannot be linear in $1/\epsilon$, but there is no known natural geometric example demonstrating this phenomenon. Indeed, there is no known lower bound, better than the trivial $\Omega(1/\epsilon)$ bound, in any concrete geometric situation. The problem, raised and addressed 20 years ago in [14], whether or not in all natural geometric scenarios of VC-dimension d, there always exists an ϵ-net of size $O(d/\epsilon)$, is still wide open. This linear (in $1/\epsilon$) upper bound has indeed been established for a few special cases, such as point objects and halfspace ranges in two and three dimensions, and point objects and disk or pseudo-disk ranges in the plane; see [14],[12], [6], [9],[16], [4] for these and some related results.

In this note we observe that the linear bound does not hold, and the smallest possible size of an ϵ-net in a very simple geometric situation (with VC-dimension 2) is not linear. Unfortunately our lower bound is only barely non-linear, providing planar geometric examples in which the minimum

*Schools of Mathematics and Computer Science, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 69978, Israel. Email: nogaa@tau.ac.il. Research supported in part by an ERC Advanced grant, by a USA-Israeli BSF grant, and by the Hermann Minkowski Minerva Center for Geometry in Tel Aviv University.

size of an ϵ-net is at least $\Omega(\frac{1}{\epsilon}w(1/\epsilon))$, where w is (a version of) the inverse Ackermann function. For convenience, we state in the following two theorems the main results demonstrating the non-linear behavior of the functions, without specifying the precise bounds obtained.

Theorem 1.1 *For every (large) positive constant C there exist n and $\epsilon > 0$ and a set X of n points in the plane, so that the smallest possible size of an ϵ-net for lines for X is larger than $C \cdot \frac{1}{\epsilon}$*

A *fat line* in the plane is the set of all points within distance μ from a line in the plane. Equivalently, this is the intersection of two half planes with parallel supporting lines. Our construction implies the following.

Theorem 1.2 *For every (large) positive constant C there exists a sequence ϵ_i of positive reals tending to zero, so that for every $\epsilon = \epsilon_i$ in the sequence and for all $n > n_0(\epsilon_i)$ there exists a set Y_n of n points in the plane, so that the smallest possible size of an ϵ-net for fat lines for Y_n is larger than $C \cdot \frac{1}{\epsilon}$*

It is not difficult to check that the VC-dimension of the range space considered in the first theorem is 2, while that of the space considered in the second is 5.

2 The proofs

We need a powerful combinatorial result of Furstenberg and Katznelson [7], [8], known as the density Hales-Jewett Theorem. For an integer $k \geq 2$, put $[k] = \{1, 2, \ldots, k\}$ and let $S = S(k, d) = [k]^d$ denote the set of all vectors of length d with coordinates in $[k]$. A *combinatorial line* is a subset $L \subset S$ so that there is a set of coordinates $I \subset [d] = \{1, 2, \ldots, d\}$, $I \neq [d]$, and values $k_i \in [k]$ for all $i \in I$ for which L is the following set of k members of S:

$$L = \{\ell_1, \ell_2, \ldots, \ell_k\}$$

where

$$\ell_j = \{(x_1, x_2, \ldots, x_d) \ : \ x_i = k_i \text{ for all } i \in I \text{ and } x_i = j \text{ for all } i \in [d] \setminus I\}.$$

Thus a combinatorial line is a set of k vectors all having some fixed values in the coordinates in I, where the jth vector has the value j in all other coordinates. In this notation, the Furstenberg-Katznelson theorem is the following.

Theorem 2.1 ([8]) *For any fixed integer k and any fixed $\delta > 0$ there exists an integer $d_0 = d_0(k, \delta)$ so that for any $d \geq d_0$, any set Y of at least $\delta k^d = \delta |S(k, d)|$ members of $S(k, d)$ contains a combinatorial line.*

We will also use the following simple lemma.

Lemma 2.2 *For every positive integer d there are d vectors $v_1, v_2, \ldots, v_d$ in the plane so that for every two nontrivial sequences of integers $(k_1, k_2, \ldots, k_d)$ and $(k'_1, k'_2, \ldots, k'_d)$, with $|k_i|, |k'_i| < k$ for all i, the two vectors $\sum_i k_i v_i$ and $\sum_i k'_i v_i$ have the same direction if and only if $(k_1, k_2, \ldots, k_d)$ and $(k'_1, k'_2, \ldots, k'_d)$ have the same direction (that is, one is a multiple of the other). Moreover, there are such vectors v_i in which all coordinates are integers of absolute value at most $(2k-1)^{2d}$.*

Proof: We show that if $v_i = (x_i, y_i)$ and each of the $2d$ numbers x_i, y_i is chosen randomly, uniformly and independently among the set of integers of absolute value at most $(2k-1)^{2d}$, then with positive probability the vectors obtained satisfy the desired properties. To prove this is indeed the case, fix two sequences $(k_1, k_2, \ldots, k_d)$ and $(k'_1, k'_2, \ldots, k'_d)$, with $|k_i|, |k'_i| < k$ and assume they are not proportional. The two vectors $\sum_i k_i(x_i, y_i)$ and $\sum_i k'_i(x_i, y_i)$ have the same direction iff

$$(\sum_i k_i x_i)(\sum_i k'_i y_i) = (\sum_i k_i y_i)(\sum_i k'_i x_i),$$

that is, iff

$$\sum_{i,j} (k_i k'_j - k_j k'_i) x_i y_j = 0.$$

As the two vectors $(k_1, k_2, \ldots, k_d)$ and $(k'_1, k'_2, \ldots, k'_d)$ are not proportional, the polynomial in the left hand side of the last equality is nontrivial and has degree 2. It thus follows, by the Schwartz-Zippel Lemma ([19], [21]) that the probability it vanishes in the random assignment to the variables x_i, y_i does not exceed $\frac{2}{2\cdot(2k-1)^{2d}+1} < \frac{1}{(2k-1)^{2d}}$. Since there are less than $(2k-1)^{2d}$ choices for the two sequences $(k_1, k_2, \ldots, k_d)$ and $(k'_1, k'_2, \ldots, k'_d)$ it follows that with positive probability none of the relevant polynomials vanishes, completing the proof. □

Proof of Theorem 1.1: Given a large positive constant C, fix an integer k satisfying $k > 2C$, let $d = d_0(k, 1/2)$ be as in Theorem 2.1 and define $n = k^d$, $\epsilon = \frac{k}{k^d}$. Let $v_1, v_2, \ldots, v_d$ be d vectors in R^2 satisfying the assertion of Lemma 2.2.

Let X be the following set of k^d points in the plane.

$$X = \{m_1 v_1 + m_2 v_2 + \ldots + m_d v_d : \ 1 \le m_i \le k \text{ for all } i\}.$$

Note that for every combinatorial line L in $S(k, d)$, the set of k points

$$\{m_1 v_1 + m_2 v_2 + \ldots m_d v_d : \ (m_1, m_2, \ldots, m_d) \in L\}$$

lies on a (geometric) line containing exactly k points of X. Indeed, if the combinatorial line is determined by the set of coordinates $I \subset [d]$, then the direction of the corresponding geometric line is $\sum_{i \in [d]-I} v_i$, and the generic choice of the vectors v_i ensures that this geometric line does not contain any additional points of X. Indeed, if $m_1 v_1 + m_2 v_2 + \ldots + m_d v_d$ is one of the points of the line, then for any other point $m'_1 v_1 + m'_2 v_2 + \ldots + m'_d v_d$ on it, the difference $(m_1 - m'_1)v_1 + (m_2 - m'_2)v_2 + \ldots + (m_d - m'_d)v_d$ must have the same direction as $\sum_{i \in [d]-I} v_i$. This implies, by the choice of the vectors v_i, that the vector $(m_1 - m'_1, m_2 - m'_2, \ldots, m_d - m'_d)$ is proportional to the characteristic vector of the set $[d] - I$, implying that indeed the vector $(m'_1, m'_2, \ldots, m'_d)$ also lies in the combinatorial line L and showing that indeed there are exactly k points of X on the corresponding geometric line.

It thus follows, by Theorem 2.1 and the choice of d, that any set of half the points of X fully contains one of these lines and thus its complement is not an ϵ-net for lines for the set X, by the definition of ϵ. Therefore, the smallest possible size of such an ϵ-net is bigger than

$$\frac{1}{2} k^d = \frac{k}{2} \frac{1}{\epsilon} > C \cdot \frac{1}{\epsilon}.$$

This completes the proof of Theorem 1.1. □

Proof of Theorem 1.2: The construction is a simple modification of the previous one. Given C, pick an integer $k > 3C$. Let $d_0 = d_0(k, 1/2)$ be as in Theorem 2.1. For each $d \geq d_0$ define $\epsilon = \epsilon(d) = 0.9\frac{k}{k^d}$ and let X be a set of k^d points in the plane defined as in the previous proof. Thus, any ϵ-net for lines for X contains at least $\frac{1}{2}k^d$ points of X. For each $n > 20 \cdot k^d$ let Y_n be a set of n points obtained from X by replacing each point x of X by a set S_x of either $\lfloor n/|X| \rfloor$ or $\lceil n/|X| \rceil$ points, all very close to x. The points in each such set S_x are chosen sufficiently close to x to ensure that for every collection of k sets that replace the points corresponding to those of a combinatorial line, there is a fat line containing all the points in these sets, and no other points of Y_n. Any subset of less than $\frac{1}{2}k^d$ of the points in Y_n must completely miss at least half of the sets S_x, and hence, by Theorem 1.1, does not intersect at least one fat line corresponding to a combinatorial line. As each such fat line is of relative size at least

$$\frac{k\lfloor n/|X| \rfloor}{n} > 0.9\frac{k}{|X|} = \epsilon,$$

this completes the proof. □

3 Concluding remarks

- The proof of [8] is not effective, and thus provides no explicit bounds, but subsequent proofs (see [17]) do provide some (very weak) estimates, and we can thus write some (extremely slowly), explicit growing function w so that the assertions of Theorems 1.1 and 1.2 hold when C is replaced by $w(\lceil 1/\epsilon \rceil)$. Indeed, the proof in [17] gives roughly the bound $A_k(1/\delta)$ for the function $d_0(k, \delta)$ defined in Theorem 2.1, where A_k is the kth function in the Ackermann hierarchy defined recursively as follows: $A_k(1) = 2$ and $A_k(n) = A_{k-1}(A_k(n-1))$, with $A_1(n) = 2n$. Thus, the kth function is obtained by iterating the (k-1)st function, so $A_2(n)$ is the exponential function 2^n and $A_3(n)$ is the tower function. Plugging in the proof of Theorems 1.1 and 1.2 we conclude that the lower bound they provide is of the form $\Omega(\frac{1}{\epsilon}w(\lceil 1/\epsilon \rceil))$, where $w(s)$ is the minimum number k so that $k^{A_k(2)} > s$.

- The example described in Theorem 1.2 clearly implies the same lower bound for several similar range spaces, like the one in which the objects are points in the plane and the ranges are planar rectangles or triangles.

- Our bounds hold for weak ϵ-nets as well. For a finite set of points X in R^m and a (possibly infinite) family of subsets $\mathcal{F}$ of R^m, a set $Y \subset R^m$ is a weak ϵ-net for X with respect to $\mathcal{F}$ if any $F \in \mathcal{F}$ that satisfies $|F \cap X| \geq \epsilon|X|$ contains at least one point of Y. The difference between this notion and that of a (strong) ϵ-net considered in the previous sections is that here Y does not have to necessarily be a subset of X. Indeed, this makes the task of finding a small net much easier, and unlike the case of strong nets it is known that there is a function $f(\epsilon, m)$ depending only on ϵ and m so that for every finite set X (of any size) in R^m, there is a weak ϵ-net Y for X with respect to the set of all convex sets in R^m, where $|Y| \leq f(\epsilon, m)$. This was first proved in [1], see also [2] and its references for several extensions. The corresponding assertion for strong nets is easily seen to be false (and indeed the VC-dimension of the family of all convex sets is infinite, in all dimensions $m \geq 2$). The best known upper bound for the function $f(\epsilon, 2)$ is $O(\frac{1}{\epsilon^2})$, proved in [1] and until recently there was no known lower bound exceeding $\Omega(1/\epsilon)$. Such a bound (of $\Omega(\frac{1}{\epsilon}\log(\frac{1}{\epsilon}))$ was proved in [5] using an elegant technique.

The results in Theorem 1.1 and in Theorem 1.2 can be easily extended to weak nets as well. Although this provides much weaker lower bounds than the above mentioned $\Omega(\frac{1}{\epsilon}\log(\frac{1}{\epsilon}))$ bound, these bounds are still non-linear in $1/\epsilon$ and hold for a very restricted collection of convex sets: lines, fat lines, rectangles or triangles, all having bounded VC-dimension. In this respect this is stronger than the results in [5]. The proofs are almost identical to the case of strong nets. For lines one first shows that a careful choice of the vectors v_i in the proof of Theorem 1.1 ensures that the only points in the plane that lie in more than two geometric lines corresponding to combinatorial lines are the points of X. This enables us to replace any weak ϵ net Y that intersects all those geometric lines (ignoring all other lines) by a strong ϵ-net for these lines, of size at most $2|Y|$. The result thus follows from (the proof of) Theorem 1.1. The derivation of the assertion of Theorem 1.2 for weak nets from the result for lines follows by essentially repeating the arguments used in the proof of Theorem 1.2. We omit the details.

- An equivalent concise way of describing the proof of Theorem 1.1 is the following. Consider the set $Z = [k]^d$ as a subset of the Euclidean space R^d. Call a (geometric) line in this space *special* if it contains all k points of one of the combinatorial lines defined in Section 1. Put $\epsilon = \frac{k}{k^d}$. The Furstenberg-Katznelson result (Theorem 2.1) implies that if d is sufficiently large, then the minimum possible size of an ϵ-net for Z with respect to the range space consisting of all special lines is at least $\frac{1}{2}k^d$. Now map Z by a random linear transformation to the plane R^2, and note that with probability 1 no two points of Z are mapped to the same point, and every special line is mapped to a line containing exactly k points of the image of Z. (This is proved algebraically in Lemma 2.2.) Let X be the image of Z. The smallest size of an ϵ-net for X is at least $\frac{1}{2}k^d$, as the set of points of Z mapped to any such net intersects all special lines.

  The result for weak nets follows in a similar way. It is easy to check that the only points in R^d that are common to at least two special lines are the points of Z, and thus any weak net for these lines in R^d can be converted into a strong net of the same cardinality, which has to be large, by Theorem 2.1. A random projection to R^2 maps Z to a set of points X and maps each special line to a line we call a special planar line. Moreover, with probability 1 the only points of the plane that belong to more than 2 of the special planar lines are the points of X. Thus for any weak ϵ-net N for X in the plane with respect to the special planar lines, its inverse image in the set of all points belonging to at least one special line is of size at most $2|N|$ and forms a weak ϵ-net for Z with respect to the set of all special lines. This shows that if d is sufficiently large, then $2|N| \geq \frac{1}{2}k^d$, providing the required estimate.

- It may be a bit better to apply Moser numbers rather than Hales-Jewett numbers (c.f., e.g. [18] for the definition of Moser numbers).

- The problem of deciding whether or not there are natural geometric range spaces of VC-dimension d in which the minimum possible size of an ϵ-net is $\Omega(\frac{d}{\epsilon}\log(1/\epsilon))$ remains open.

Acknowledgment I would like to thank Esther Ezra and Shakhar Smorodinsky for helpful discussions.

References

[1] N. Alon, I. Bárány, Z. Füredi and D. J. Kleitman, Point selections and weak ϵ-nets for convex hulls, Combinatorics, Probability and Computing 1 (1992), 189-200.

[2] N. Alon, G. Kalai, J. Matoušek and R. Meshulam, Transversal Numbers for Hypergraphs Arising in Geometry, Advances in Applied Math. 29 (2002), 79-101.

[3] N. Alon and J. H. Spencer, *The Probabilistic Method, Third Edition*, Wiley, 2008, xv+352 pp.

[4] B. Aronov, E. Ezra and M. Sharir, Small-size Epsilon-Nets for Axis-Parallel Rectangles and Boxes Proc. STOC 09, 639-648.

[5] B. Bukh, J. Matoušek and G. Nivasch, Lower bounds for weak epsilon-nets and stair-convexity Proc. 25th ACM Symp. on Computational Geometry (SoCG 2009), pp. 1-10, 2009.

[6] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set cover, Discrete Comput. Geom. 37 (2007), 43–58.

[7] H. Furstenberg and Y. Katznelson, A density version of the Hales-Jewett theorem for k = 3, Graph Theory and Combinatorics (Cambridge, 1988), Discrete Math. 75 (1989), 227-241.

[8] H. Furstenberg and Y. Katznelson, A density version of the Hales-Jewett theorem, J. Anal. Math. 57 (1991), 64-119.

[9] S. Har-Peled, H. Kaplan, M. Sharir, and S. Smorodinsky, ϵ-nets for halfspaces revisited, manuscript, 2008

[10] D. Haussler and E. Welzl, ϵ-nets and simplex range queries, Discrete and Computational Geometry 2 (1987), 127-151.

[11] J. Komlós, J. Pach, and G. Woeginger, Almost tight bounds for epsilon nets, Discrete Comput. Geom. 7 (1992), 163–173.

[12] J. Matoušek, Reporting points in halfspaces, Comput. Geom. Theory Appl. 2 (1992), 169–186.

[13] J. Matoušek, *Lectures on Discrete Geometry*, Graduate Texts in Mathematics, 212. Springer-Verlag, New York, 2002, xvi+481 pp.

[14] J. Matoušek, R. Seidel and E. Welzl, How to net a lot with little: Small ϵ-nets for disks and halfspaces, In Proc. 6th Annu. ACM Sympos. Comput. Geom., pages 16-22, 1990. Revised version at http://kam.mff.cuni.cz/ matousek/enets3.ps.gz.

[15] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, Wiley-Interscience Series in Discrete Mathematics and Optimization. A Wiley-Interscience Publication. John Wiley and Sons, Inc., New York, 1995, xiv+354 pp.

[16] E. Pyrga and S. Ray, New existence proofs for ϵ-nets, Proc. 24th Annu. ACM Sympos. Comput. Geom. (2008), 199–207.

[17] D. H. J. Polymath, A new proof of the density Hales-Jewett theorem, preprint, available at arxiv.org/abs/0910.3926.

[18] D. H. J. Polymath, Density Hales-Jewett and Moser numbers, preprint, available at arXiv:1002.0374.

[19] J. Schwartz, Fast probabilistic algorithms for verification of polynomial identities, Journal of the ACM (1980), 701–717.

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

[21] R. Zippel, Probabilistic algorithms for sparse polynomials, Symbolic and algebraic computation (EUROSAM '79, Internat. Sympos., Marseille, 1979), pp. 216–226, Lecture Notes in Comput. Sci., 72, Springer, Berlin-New York, 1979.