Weak ε-nets for Axis-Parallel Boxes in d-Space*

Esther Ezra†

May 25, 2009

Abstract

In this note we show the existence of weak ε-nets of size $O(1/\varepsilon \log \log (1/\varepsilon))$ for point sets and axis-parallel boxes in $\mathbb{R}^d$. Our analysis uses a non-trivial variant for the recent technique of Aronov *et al.* [AES09] that yields (strong) ε-nets, whose size have the above asymptotic bound, for $d = 2, 3$.

*Work on this paper has been supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, and by an NIH grant 1P50-GM-08183-01.

†Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; *esther@cs.duke.edu*.

1 Introduction

Let P be a set of n points in $\mathbb{R}^d$, and let $0 < \varepsilon < 1$ be a parameter. An *ε-net* for P with respect to a set $\mathcal{R}$ of regions in d-space is a subset $N \subseteq P$ with the property that any region $R \in \mathcal{R}$ with $|R \cap P| \geq \varepsilon|P|$ contains an element of N. In other words, N is a hitting set for all the "heavy" ranges.

The epsilon-net theorem of Haussler and Welzl [HW87], a fundamental result in this area, asserts that, for any pair $(P, \mathcal{R})$ (also called a *range space*), and ε as above, such that $(P, \mathcal{R})$ has *finite VC-dimension*[1] d, there exists an ε-net N of size $O\left(d/\varepsilon \log\left(d/\varepsilon\right)\right)$, and that in fact a random sample of P of that size is an ε-net with constant probability. In particular, the size of N is independent of the size of P.

A major open problem in the theory of ε-nets is related to the question whether the factor $\log \frac{1}{\varepsilon}$ in the upper bound on their size is really necessary, especially for simply-shaped (low-dimensional) regions. The reader is referred to [AES09] for the discussion. The prevailing conjecture is that, at least in these geometric scenarios, there always exists an ε-net of size $O(1/\varepsilon)$ [MSW90] (which matches the best lower bound that is currently known). This "linear" upper bound has indeed been established for a few special cases, such as points and halfspaces in two and three dimensions, points and disks or pseudo-disks in the plane, and points and translates of a fixed convex polytope in 3-space; see [MSW90, Mat92b, CV07, HKSS08, PR08, 1].

When the range space $(P, \mathcal{R})$ does not have a finite VC-dimension, the bound obtained by the epsilon-net theorem does not necessarily hold. A well-studied family is the case where $\mathcal{R}$ consists of all convex subsets of $\mathbb{R}^d$ (more precisely, each set in $\mathcal{R}$ is the intersection of P with a convex region in $\mathbb{R}^d$). In this case, it can be shown that the ε-net size can be at least $(1-\varepsilon)n$ [HW87]. To overcome this difficulty, Haussler and Welzl introduced the concept of *weak ε-nets*. In this case the net is a set N satisfying the above conditions, except the requirement that N be a subset of P. In some cases ,the fact that N can be chosen anywhere in space aids us to construct weak ε-nets of relatively small size. Returning to the case of points and convex subsets of $\mathbb{R}^d$, this crucial relaxation of the definition enables us to construct weak ε-nets whose size depends only in ε (and not in n), in this case the ε-net size is $O(\varepsilon^{-d} \operatorname{polylog}(1/\varepsilon))$ [CEGGSW95, MW04] (which is an improvement of the result in [ABFK92]); see also a recent result by Alon *et al.* [AKNSS08] for the case where all the points in P are in convex position, and [MR08] for a connection between weak ε-nets with respect to convex sets and "strong" ε-nets (that is, when N is a subset of P) with respect to other set systems with finite VC-dimension. Concerning the lower bound, Bukh *et al.* [BMN09] have very recently shown that there are range spaces of points and convex sets in d-space that admit a weak ε-net of size $\Omega(1/\varepsilon \log^{d-1}(1/\varepsilon))$, thus showing the first super-linear lower bound construction (nevertheless, when the regions are of constant description complexity, the best lower bound is still linear).

Our result. In this note we consider the case of points and axis-parallel boxes in d-space, for any $d \geq 4$, and show that this range space admits a weak ε-net of size $O(1/\varepsilon \log\log(1/\varepsilon))$ (with a constant of proportionality that depends on d), thus significantly improving the currently known bound $O(1/\varepsilon \log(1/\varepsilon))$ (simply obtained by the epsilon-net theorem). Our technique is a non-trivial variant of the technique presented by Aronov *et al.* [AES09] for obtaining (strong) ε-nets for

[1]Informally, it suffices to require that the number of distinct sets $R \cap P$, for $R \in \mathcal{R}$, is polynomial in $|P|$. This scenario includes the case where the regions in $\mathcal{R}$ are simply-shaped, formally, they are assumed to have *constant descriptive complexity*, meaning that they are semi-algebraic sets defined in terms of a constant number of polynomial equations and inequalities of constant maximum degree.

range spaces of this kind, whose size has a similar asymptotic bound, for $d = 2, 3$. The technique in [AES09] fails to produce a small-size (strong) ε-nets when $d \geq 4$, which arises a challenging open problem: Does the bound $O(1/\varepsilon \log\log(1/\varepsilon))$ holds in any dimension? In this note we relax this problem, and construct a weak ε-net of this size, thus resolving the aforementioned problem in this particular scenario.

In our analysis, we construct an ε-net which is a combination of the construction in [AES09] (and thus contains a subset of P) and a set of "artificial points" (not from P) that constitute the "weakness" of the net.

2 Preliminaries

For the sake of completeness, and in order to simplify the presentation in Section 3, we first present an overview of the technique of Aronov *et al.* [AES09] to obtain a (strong) ε-net for points and axis-parallel rectangles in the plane. The essence of their approach is *oversampling*, that is, sampling a slightly larger set of points of the input, in order to guarantee that most of the "heavy" rectangles are indeed "pierced"' by the net.

The sampling technique. Let P be a set of n points in the plane. Put $r := 4/\varepsilon$, $s := cr \log\log r$, where $c > 2$ is a sufficiently large constant. We first draw a random sample R so that each point $p \in P$ is chosen independently to be included in R with probability $\pi := s/n$; thus the expected size of R is[2] s. We make R part of the ε-net to be constructed. We then need to handle only axis-parallel rectangles which contain at least n/r points, but are *R-empty*, i.e., (axis-parallel) rectangles which do not contain any point of R. To pierce every such rectangle, we form the subset $\mathcal{M}$ of *maximal R-empty rectangles*, that is, rectangles, each of which is determined by at most four points of R, one on each of its sides (in case it is bounded there). Any other R-empty rectangle Q is contained in one of them (by simply pushing each of its sides until it touches a point of R or extends to $\pm\infty$). By the standard ε-net theory of [HW87], with high probability, each rectangle of $\mathcal{M}$ contains at most $O\left(\frac{n}{s}\log s\right)$ points of P. Moreover, the expected number of points of P in such a rectangle is $O(n/s)$. Since $s \gg r$, most rectangles of $\mathcal{M}$ contain fewer than $\varepsilon n = n/r$ points of P, so an R-empty rectangle Q with at least n/r points will not fit into any of them, and we can simply ignore them. For each of the relatively few "heavy" rectangles M of $\mathcal{M}$, we apply the resampling technique of [CF90, CV07], also referred to as a *repair sampling step*, as follows.

Define the *weight factor* t_M of M to be $s|M \cap P|/n$. Rectangles M with $t_M < s/r = c \log\log r$ can be ignored, because they contain fewer than n/r points of P, so no rectangle Q, as above, can be completely contained in one of them. By the standard ε-net theory [HW87], for each $M \in \mathcal{M}$ with $t_M \geq c \log\log r$, there exists a subset $N_M \subseteq M \cap P$ of size $c' t_M \log t_M$ that forms a $(1/t_M)$-net for $M \cap P$, where $c' > 0$ is an absolute constant. The final ε-net N is the union of R with the sets N_M, over all the "heavy" rectangles $M \in \mathcal{M}$ (i.e., rectangles with $t_M \geq c \log\log r$). The (fairly easy) analysis in [AES09] shows that this is indeed an ε-net.

The problem decomposition. As shown in [AES09], the number of maximal R-empty rectangles can be $\Theta(s^2)$ in the worst case, which leads, using the above construction literally, to a $\Theta\left(\frac{1}{\varepsilon^2}\right)$ bound on the expected size of the ε-net in the worst case, which is much too large. To overcome

[2] From technical reasons described in [AES09], it is crucial to choose this sampling model.

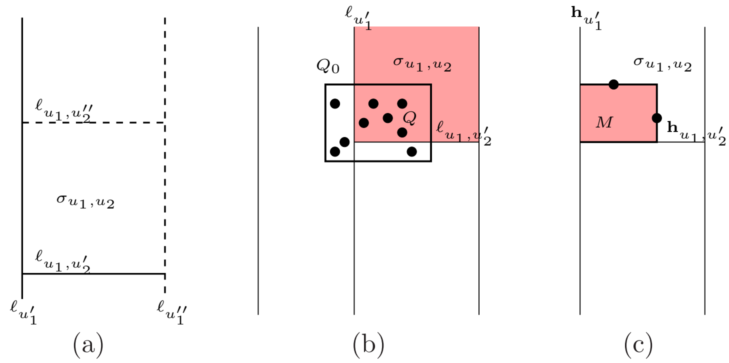

Figure 1: (a) The "quadrant" σ_{u_1,u_2} is defined by the line splitters $\ell_{u_1'}$, $\ell_{u_1,u_2'}$, but it is also bounded by ancestor splitters $\ell_{u_1''}$ and $\ell_{u_1,u_2''}$. (b) The rectangle Q is anchored at the (apex of the) quadrant σ_{u_1,u_2}. (c) An anchored rectangle M that is determined by a pair of points.

this difficulty, we observe that it is sufficient to consider much fewer maximal R-empty rectangles. To do so, we construct a two-level range-tree $\mathcal{T}$, over the points of P (see, e.g., [dBCKO08]), where the points are sorted by their x-coordinates in the primary tree, and by their y-coordinates in each secondary tree[3]. We associate with each node u_1 of the primary tree the subset P_{u_1} of points that it represents, and a secondary (y-sorted) tree $\mathcal{T}_{u_1}$ on P_{u_1}. Similarly, with each node u_2 of a secondary tree $\mathcal{T}_{u_1}$ we associate the corresponding subset P_{u_1,u_2} of P_{u_1}. We construct each of the two levels of $\mathcal{T}$ down to nodes for which the size of their associated subset is between n/r and $n/(4r)$. Each of the primary and secondary trees has at most $2 + \log r$ levels, and the total number of nodes in the range-tree $\mathcal{T}$ is $O(r \log r)$.

Each non-leaf node u_1 of the primary tree stores a vertical line ℓ_{u_1} that evenly splits the points in $P(u_1)$ into two subsets stored at the children of u_1, and each non-leaf node u_2 of a secondary tree $\mathcal{T}_{u_1}$ stores a horizontal line "splitter" ℓ_{u_1,u_2} with similar properties. For each such secondary node u_2 of a tree $\mathcal{T}_{u_1}$, the lines $\ell_{u_1'}$ and $\ell_{u_1,u_2'}$, where u_1' is the parent of u_1 in $\mathcal{T}$ and u_2' is the parent of u_2 in $\mathcal{T}_{u_1}$, define a quadrant σ_{u_1,u_2}, which is the intersection of two halfplanes bounded by $\ell_{u_1'}$ and $\ell_{u_1,u_2'}$ and containing P_{u_1,u_2}. As discussed in [AES09], we only need to consider pairs (u_1, u_2) of vertices, each of which has a parent in its respective tree. Thus each of the two halfplanes are proper, and σ_{u_1,u_2} is a non-degenerate quadrant. More precisely, it is more accurate to regard σ_{u_1,u_2} as a box, or a clipped quadrant, bounded on the other side also by lines associated with ancestors of u_1 and u_2 (if exist); see Figure 1(a).

Given an axis-parallel (R-empty) rectangle Q_0 containing at least εn, we first locate the highest node u_1' in $\mathcal{T}$, so that the line $\ell_{u_1'}$ meets Q_0, thus splitting Q_0 into two parts, one of which, call it Q_1, contains at least $\varepsilon n/2$ points of P. Let u_1 be the corresponding child of u_1' so that H_{u_1} contains Q_1. We next locate the highest node u_2' in $\mathcal{T}_{u_1}$, such that $\ell_{u_1,u_2'}$ meets Q_1. We focus on the portion Q of Q_1 that contains at least $\varepsilon n/4 = n/r$ points, and denote by u_2 the child of u_2' whose corresponding quadrant σ_{u_1,u_2} contains Q. By construction, Q is *anchored* at the resulting quadrant $\sigma := \sigma_{u_1,u_2}$, in the sense that the apex o of σ is a vertex of Q, moreover, Q behaves as a quadrant within σ that is oppositely oriented to σ. See Figure 1(b).

By the above considerations, it is sufficient to consider all maximal anchored R-empty rectangles

[3]In fact, the actual approach in [AES09] for the two-dimensional case is somewhat simpler, and uses just a binary-tree-like hierarchy of vertical strips, whereas a range-tree decomposition is used only for the three-dimensional problem. Nevertheless, we choose to present this variant, since the analysis given in this paper exploits a structure of this kind.

M within σ, and apply the repair sampling step in each of them (it is then easy to verify that Q is contained in at least one such rectangle M). Let $\mathcal{M}_{u_1,u_2}$ be the set of all such rectangles. Since each of these rectangles behaves as a quadrant inside σ (and is determined by at most two points of $R_{u_1,u_2} := R \cap \sigma$), the number of such empty quadrants is only $O(1 + r_{u_1,u_2})$ [BSTY98, KRSV07], where $r_{u_1,u_2} := |R_{u_1,u_2}|$; see Figure 1(c). Summing these bounds over all nodes u_2 of all secondary trees $\mathcal{T}_{u_1}$, appearing in a fixed pair of levels i_1 (primary) and i_2 (secondary), it is easy to verify that the resulting bound is $O(|R| + r)$, for a total of $O(|R| \log^2 r + r \log r)$ over all pairs of levels (i_1, i_2), thus yielding a bound on the overall size of the sets $\mathcal{M}_{u_1,u_2}$ (see [AES09] for further details).

The Exponential Decay Lemma. We next bound the expected size of the final ε-net N. Define $\mathrm{CT}(R)$ to be the union of all the collections $\mathcal{M}_{u_1,u_2}$, over all nodes u_2 of all secondary trees $\mathcal{T}_{u_1}$, appearing in a fixed pair of levels (i_1, i_2). For a positive parameter t, let $\mathrm{CT}_t(R)$ be the subset of $\mathrm{CT}(R)$ consisting of those rectangles M with $t_M \geq t$. A crucial ingredient of the analysis in [AES09] is the so-called Exponential Decay Lemma, which implies (see also [AMS98]):

$$\mathbf{Exp}\left\{|\mathrm{CT}_t(R)|\right\} = O\left(2^{-t}\, \mathbf{Exp}\left\{|\mathrm{CT}(R')|\right\}\right),$$

where $\mathbf{Exp}(\cdot)$ denotes expectation, and R' is another random sample of P, where each point $p \in P$ is now chosen, independently, to belong to R' with probability $\pi' := \pi/t$. We apply the lemma with $t = c \log\log r$, so $\pi' = \pi/t = r/n$. By the above considerations, $\mathbf{Exp}\left\{|\mathrm{CT}(R')|\right\} = O(r)$. We thus have

$$\mathbf{Exp}\left\{|\mathrm{CT}_t(R)|\right\} = O\left(2^{-t}\, \mathbf{Exp}\left\{|\mathrm{CT}(R')|\right\}\right) = O\left(r 2^{-c\log\log r}\right) = O\left(r/\log^c r\right).$$

In other words, the expected number of "heavy" rectangles of $\mathrm{CT}(R)$ is only sublinear in r. Moreover, the analysis in [AES09] shows:

$$\mathbf{Exp}\left\{ \sum_{(u_1,u_2) \text{ at levels } (i_1,i_2)} \sum_{\substack{M \in \mathcal{M}_{u_1,u_2} \\ t_M \geq c\log\log r}} t_M \log t_M \right\} = O\left(\frac{r \log\log r \log\log\log r}{\log^c r}\right).$$

Repeating the analysis for each of the $O\left(\log^2 r\right)$ pairs (i_1, i_2), the expectation of the above sum is $o(r)$, provided $c > 2$, as we indeed assume; thus (recall that N is the final ε-net)

$$\mathbf{Exp}\{|N|\} = \mathbf{Exp}\{|R|\} + o(r) = O(r \log\log r).$$

One of the major properties of the analysis above is to have a (nearly) linear bound on the expected size of $\mathrm{CT}(R')$. By the properties of this construction, any bound of the form $O(r \operatorname{polylog} r)$ will yield a sublinear bound on the number of points sampled at the repair sampling step, given that c is sufficiently large. We will return to this property in Section 3.

The natural extension of this technique to three dimensions yields an ε-net whose size has a similar asymptotic bound. This follows due to the fact that there is only a linear number of maximal anchored R-empty boxes M in each orthant σ of the decomposition, thus yielding a linear bound on $\mathbf{Exp}\{|\mathrm{CT}(R')|\}$. This property is, however, violated in dimensions $d \geq 4$. In this case, the bound can be $\Theta\left(|R \cap \sigma|^{\lfloor d/2 \rfloor}\right)$ in the worst case, (see [BSTY98, KRSV07]), which is at least quadratic for $d \geq 4$, and thus yields an ε-net of expected size $\Theta\left(r^{\lfloor d/2 \rfloor}\right)$ in the worst case, using the technique outlined above literally.

3 The Construction

We now present our construction for the weak ε-net. Let P be a set of n points in $\mathbb{R}^d$. Put $r := 2^d/\varepsilon$ and $s := cr \log\log r$, for some constant $c > 0$ that we fix shortly. We first choose each point of P independently with probability s/n, thereby producing a sample R whose expected size is s, and make it part of the output.

We next construct a d-level range-tree $\mathcal{T}$, over the points of P, where the points are sorted by their x_i-coordinates at each i-ary tree, for $i = 1, \ldots, d$. We construct each of the d levels of $\mathcal{T}$ down to nodes for which the size of their associated subset is between $n/(2^d r)$ and n/r. Clearly, each of the i-ary trees, for $i = 1, \ldots, d$, has at most $d + \log r$ levels, and the total number of nodes in the range-tree $\mathcal{T}$ is $O(r \log^{d-1} r)$.

We use a similar sampling approach as in Section 2, but the repair sampling step is applied only at those nodes u_d of the d-ary trees that lie at the low levels (starting from the leaves) of these trees, and thus their respective orthants σ contain a relatively small number of points of R. This guarantees, as shown by the analysis, that the overall number of points sampled at the second sampling step is indeed small. Nodes that lie at the higher levels of the d-ary trees (in which case $\sigma \cap R$ may be large) are handled differently, and we choose for each of them the apex o of the orthant σ associated with them (if exits) to be part of the final ε-net. This guarantees that all boxes anchored at σ are stabbed by o. The last modification imposes the "weakness" of the ε-net, that is, the apices of the orthants σ are not necessarily part of the point set P, however, this is exactly the case where the analysis of [AES09] fails (for dimensions $d \geq 4$), as the number of points added to the output at the second sampling step (for nodes at the higher levels) is much too large, whereas the construction presented here skips that step for this set of nodes—see below.

Analogously to the construction in [AES09] overviewed in Section 2 and using the notation there, each non-leaf node u_i of the i-ary tree stores a halving hyperplane $\mathbf{h}_{u_1,\ldots,u_i}$, which we denote with a slight abuse of notation by $\mathbf{h}_{u_i}$, orthogonal to the x_i-axis, which evenly splits the subset of points of P (and R) associated with u_i, for $i = 1, \ldots, d$. Next, let B_0 be an axis-parallel box containing at least εn points of P. Using the considerations in [AES09] (and Section 2), we query $\mathcal{T}$ with B_0 and obtain a sequence of d vertices $u'_1, \ldots, u'_d$, each of which is associated with a halving hyperplane $\mathbf{h}_{u'_i}, \ldots, \mathbf{h}_{u'_d}$ orthogonal to the x_i-axis, respectively, such that one of the orthants σ induced by these hyperplanes captures at least $\varepsilon n/2^d$ points of B_0, that is, $|B_0 \cap \sigma| \geq \varepsilon n/2^d$. Let u_i be the child of u'_i at the i-ary tree, whose corresponding halfspace H_{u_i} (bounded by $\mathbf{h}_{u'_i}$) contains σ, for $i = 1, \ldots, d$. We thus have $\sigma = \sigma_{u_1,\ldots,u_d} = \bigcap_{i=1,\ldots,d} H_{u_i}$. Put $B := B_0 \cap \sigma$.

By construction, B is anchored at the resulting orthant σ. Moreover, the apex o of σ stabs any box that is anchored at σ. We thus output all apices o of the corresponding orthants σ associated with nodes $u = u_d$ of the d-ary trees $\mathcal{T}^d$, for which the level of u at $\mathcal{T}^d$ is at least $c' \log\log r$ (beginning from the leaves of $\mathcal{T}^d$), where $c' > 0$ is a constant that we fix shortly. In other words, we ignore, for the time being, the lower $c' \log\log r$ levels of the d-ary trees, and consider at that step only the higher ones. Clearly, the overall number of apices that we collect in each d-ary tree is

$$O\left(\left\lfloor 2^{j - c' \log\log r} \right\rfloor\right) = O\left(\left\lfloor \frac{2^j}{\log^{c'} r} \right\rfloor\right),$$

where $1 \leq j \leq \log r$ is the number of levels at that tree. It is easy to verify that the sum of these bounds, over all d-ary trees that are associated with a fixed $(d-1)$-ary tree is $O(r \log r / \log^{c'} r)$, and, using induction on d, the overall sum of these bounds, over all d-ary trees of $\mathcal{T}$ is

$$O\left(\frac{r\log^{d-1} r}{\log^{c'} r}\right),$$

which is $o(r)$ if we choose $c' > d-1$.

When u lies at level at most $c' \log\log r$, we apply the repair sampling step. Put $R_u := R \cap \sigma$, where σ is the orthant associated with u, and $r_u := |R_u|$. Let $\mathcal{M}_u$ be the set of all maximal anchored R-empty (i.e., R_u-empty) axis-parallel boxes contained in the orthant σ; each of these boxes $M \in \mathcal{M}_u$ behaves as an octant inside σ, and is thus determined by at most d points of R_u, each lying on a distinct facet of M. As shown in [KRSV07], the number of these empty boxes is $O\left(r_u^{\lfloor d/2 \rfloor}\right)$. In particular, $\mathbf{Exp}\{r_u\} = O\left(\log^{c'} r\right)$, since it lies at level at most $c' \log\log r$. In order to bound the expected number of empty boxes in $\mathcal{M}_u$, we argue (similarly to the analysis in [AES09]), as follows. Put $z := \mathbf{Exp}\{r_u\}$. By Chernoff's bound (see, e.g., [AS92]),

$$\Pr\left\{r_u \geq \xi z\right\} \leq e^{-(\xi-1)^2 z/3},$$

for any $\xi > 1$. Hence,

$$\begin{aligned}\mathbf{Exp}\left\{r_u^{\lfloor d/2 \rfloor}\right\} &\leq (2z)^{\lfloor d/2 \rfloor} + \sum_{j \geq 2} \Pr\left\{jz \leq r_u \leq (j+1)z\right\} ((j+1)z)^{\lfloor d/2 \rfloor} \\ &\leq z^{\lfloor d/2 \rfloor} \cdot \left(2^{\lfloor d/2 \rfloor} + \sum_{j \geq 2} (j+1)^{\lfloor d/2 \rfloor} e^{-(j-1)^2 z/3}\right) = O\left(z^{\lfloor d/2 \rfloor}\right).\end{aligned}$$

In particular, substituting $z = O\left(\log^{c'} r\right)$, we obtain $\mathbf{Exp}\left\{r_u^{\lfloor d/2 \rfloor}\right\} = O\left(\log^{c'\lfloor d/2 \rfloor} r\right)$.

We next traverse all the "heavy" boxes $M \in \mathcal{M}_u$ with weight factor $t_M \geq c \log\log r$, and take a $(1/t_M)$-net N_M, for the set $P \cap M$, of size $O(t_M \log t_M)$ (whose existence is guaranteed by [HW87]), and report all the resulting nets N_M as part of the output, over all the heavy boxes $M \in \mathcal{M}_u$ and vertices u as above. As discussed in Section 2, every anchored box B that is contained in M and containing at least n/r points of P is hit by N_M.

We now bound the overall expected size of the nets N_M. Analogously to Section 2, we define (i) $\mathrm{CT}(R)$ to be the union of all the collections $\mathcal{M}_u$, over all nodes u of all d-ary trees $\mathcal{T}^d$, appearing in a fixed d-tuple of levels $(i_1, i_2, \ldots, i_d)$, with $i_d < c' \log\log r$, and (ii) $\mathrm{CT}_t(R)$ to be the subset of $\mathrm{CT}(R)$ consisting of those boxes M with $t_M \geq t$, for any parameter t. By the Exponential Decay Lemma we have

$$\mathbf{Exp}\left\{|\mathrm{CT}_t(R)|\right\} = O\left(2^{-t}\,\mathbf{Exp}\left\{|\mathrm{CT}(R')|\right\}\right),$$

where R' is defined as in Section 2. In particular, when $t \geq c \log\log r$ we have

$$\begin{aligned}\mathbf{Exp}\left\{|\mathrm{CT}(R')|\right\} &= \mathbf{Exp}\Bigg\{\sum_{\substack{u \text{ at levels } (i_1,\ldots,i_d) \\ i_d < c' \log\log r}} |\mathcal{M}_u|\Bigg\} \\ &= \mathbf{Exp}\Bigg\{\sum_{\substack{u \text{ at levels } (i_1,\ldots,i_d) \\ i_d < c' \log\log r}} O\left(r_u^{\lfloor d/2 \rfloor}\right)\Bigg\} = \sum_{\substack{u \text{ at levels } (i_1,\ldots,i_d) \\ i_d < c' \log\log r}} \mathbf{Exp}\Big\{O\left(r_u^{\lfloor d/2 \rfloor}\right)\Big\} \\ &= \sum_{\substack{u \text{ at levels } (i_1,\ldots,i_d) \\ i_d < c' \log\log r}} O\left(\log^{c'\lfloor d/2 \rfloor} r\right) = O\left(r \cdot \log^{c'(\lfloor d/2 \rfloor - 1)} r\right),\end{aligned}$$

where the last bound follows from the fact that any fixed d-tuple $(i_1, \ldots, i_d)$ induces a decomposition of space into $O(r/2^{i_d})$ cells (or truncated orthants σ), each of which contributes $O\left(2^{i_d \lfloor d/2 \rfloor}\right)$ maximal anchored empty boxes, according to the preceding arguments. Thus the overall contribution is $O\left(r 2^{i_d(\lfloor d/2 \rfloor - 1)}\right)$, which is bounded by the expression that we have obtained when $i_d < c' \log\log r$.

In other words, the expected number of maximal anchored empty boxes in a fixed d-tuple of levels $(i_1, i_2, \ldots, i_d)$ as above is only $O(r \operatorname{polylog} r)$, which, according to the analysis in [AES09] and the discussion in Section 2, is exactly the case that yields a sublinear bound on the overall size of the various nets N_M. Specifically, using similar considerations as in [AES09]:

$$
\begin{aligned}
&\mathbf{Exp}\Bigg\{ \sum_{\substack{u \text{ at levels} \\ (i_1,\ldots,i_d)}} \sum_{\substack{M \in \mathcal{M}_u \\ t_M \geq c \log\log r \\ i_d < c' \log\log r}} t_M \log t_M \Bigg\} \\
&= \mathbf{Exp}\Bigg\{ \sum_{j \geq t} \sum_{\substack{M \in \mathrm{CT}(R) \\ t_M = j}} j \log j \Bigg\} \\
&= \mathbf{Exp}\Bigg\{ \sum_{j \geq t} j \log j \cdot \left(|\mathrm{CT}_j(R)| - |\mathrm{CT}_{j+1}(R)| \right) \Bigg\} \\
&= \mathbf{Exp}\Bigg\{ t \log t \cdot |\mathrm{CT}_t(R)| + \sum_{j > t} \big(j \log j - (j-1)\log(j-1) \big) |\mathrm{CT}_j(R)| \Bigg\} \\
&= O\left(\frac{r \cdot \log^{c'(\lfloor d/2 \rfloor - 1)} r}{\log^c r} (t \log t) + \sum_{j > t} \frac{r \cdot \log^{c'(\lfloor d/2 \rfloor - 1)} r}{2^j} \log j \right) \\
&= O\left(\frac{r \cdot \log^{c'(\lfloor d/2 \rfloor - 1)} r \cdot t \log t}{\log^c r} \right) \\
&= O\left(\frac{r \cdot \log^{c'(\lfloor d/2 \rfloor - 1)} r \cdot \log\log r \log\log\log r}{\log^c r} \right).
\end{aligned}
$$

Repeating the analysis for each of the $O(\log^d r)$ d-tuples $(i_1, \ldots, i_d)$, we obtain that the expectation of the above sum is $o(r)$, provided $c > d + c'(\lfloor d/2 \rfloor - 1)$.

Summing each of the three bounds, obtained at the preliminary sample, the secondary sample, and the set of the orthant apices associated with d-ary nodes at levels $\geq c' \log\log r$, it follows that the expected size of the ε-net is $O(r \log\log r)$, as asserted.

Constructing the ε-net. Constructing an ε-net of this size is an easy extension of the algorithm given at [AES09]. We first construct a d-level range-tree $\mathcal{T}$ over the points of P, stopping at nodes u (at each level of the tree) for which $n/(2^d r) < |P_u| < n/r$. We next set the pair of constants c, c' to satisfy the condition stated above. We report, as part of the final ε-net, the preliminary random sample R (drawn according to the model assumed above), and all apices o of the corresponding orthants σ associated with nodes u of the d-ary trees $\mathcal{T}^d$, whose level at $\mathcal{T}^d$ is at least $c' \log\log r$. We next consider all the remaining nodes u (at level at most $c' \log\log r$ at $\mathcal{T}^d$), and process them in a similar manner as in [AES09]. Specifically, for each such node u, we enumerate all maximal anchored R-empty octants M in σ_u in $O\left(r_u^{\lfloor d/2 \rfloor} \log^{d-1} r\right)$ time, using the technique of Kaplan *et al.* [KRSV07].

Using the range-tree, we compute, for each of these octants M, its weight factor t_M, and collect those M whose weight factor exceeds $c \log \log r$. We then report, for each of the heavy octants, the set $P \cap M$, and construct a $(1/t_M)$-net of size $O(t_M \log t_M)$ for that set. Omitting any further details, we obtain that the overall expected running time is $O(n \log^{d-1} n)$.

As observed in [AES09], and using the (fairly easy) technique there, the (expected) running time can be slightly improved to $O(n \log^d r)$.

References

[AMS98] P. K. Agarwal, J. Matoušek, and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.*, 27:491–505, 1998.

[ABFK92] N. Alon, I. Barany, Z. Fredi, and D. Kleitman. Point selections and weak ε-nets for convex hulls. *Combin. Probab. Comput.*, 1(3):103–112, 1992.

[AKNSS08] N. Alon, H. Kaplan, G. Nivasch, M. Sharir, and S. Smorodinsky, Weak ε-nets and interval chains, *J. ACM*, 55, Article 28, 2008.

[AS92] N. Alon and J. Spencer, *The Probabilistic Method*, Wiley-Interscience, New York, 1992.

[AES09] B. Aronov, E. Ezra, and M. Sharir, Small-size ε-nets for axis-parallel rectangles and boxes. In *Proc. 41th Annu. ACM Symp. Theory Comput.*, to appear.

[BSTY98] J.D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discrete Comput. Geom.*, 19:485–519, 1998.

[BMN09] B. Bukh, J. Matoušek, and G. Nivasch, Lower bounds for weak ε-nets and stair-convexity, In *Proc. 25th Annu ACM Symp. Comput. Geom.*, to appear.

[CEGGSW95] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, M. Sharir, and E. Welzl. Improved bounds on weak ε-nets for convex sets. *Discrete Comput. Geom.*, 13:1–15, 1995.

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

[CV07] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set cover, *Discrete Comput. Geom.*, 37:43–58, 2007.

[dBCKO08] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd rev. ed., Springer-Verlag, Berlin, 2008.

[HKSS08] S. Har-Peled, H. Kaplan, M. Sharir, and S. Smorodinsky, ε-nets for halfspaces revisited, manuscript, 2008.

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

[KRSV07] H. Kaplan, N. Rubin, M. Sharir, and E. Verbin, Efficient colored orthogonal range counting, *SIAM J. Comput.*, 38:982–1011, 2008.

[1] S. Lauen. Geometric set cover and hitting sets for polytopes in $\mathbb{R}^3$. *Proc. 25th Int. Sympos. Theoret. Aspects Comput. Sci.*, pages 479–490, 2008.

[Mat92b] J. Matoušek, Reporting points in halfspaces, *Comput. Geom. Theory Appl.*, 2:169–186, 1992.

[MSW90] 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` .

[MW04] J. Matoušek and U. Wagner, New constructions of weak ε-Nets. *Discrete Comput. Geom.*, 32(2):195–206, 2004.

[MR08] N. H. Mustafa and S. Ray, Weak ε-nets have a basis of size $O(1/\varepsilon \log(1/\varepsilon))$, *Comput. Geom. Theory Appl.*, 40(1):84–91, 2008.

[PR08] E. Pyrga and S. Ray, New existence proofs for ε-nets, *Proc. 24th Annu. ACM Sympos. Comput. Geom.*, pages 199–207, 2008.