

SMALL-SIZE ε -NETS FOR AXIS-PARALLEL RECTANGLES AND BOXES*

BORIS ARONOV[†], ESTHER EZRA[‡], AND MICHA SHARIR[§]

Abstract. We show the existence of ε -nets of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$ for planar point sets and axis-parallel rectangular ranges. The same bound holds for points in the plane and “fat” triangular ranges and for point sets in \mathbf{R}^3 and axis-parallel boxes; these are the first known nontrivial bounds for these range spaces. Our technique also yields improved bounds on the size of ε -nets in the more general context considered by Clarkson and Varadarajan. For example, we show the existence of ε -nets of size $O\left(\frac{1}{\varepsilon} \log \log \log \frac{1}{\varepsilon}\right)$ for the dual range space of “fat” regions and planar point sets (where the regions are the ground objects and the ranges are subsets stabbed by points). Plugging our bounds into the technique of Brönnimann and Goodrich or of Even, Rawitz, and Shahar, we obtain improved approximation factors (computable in expected polynomial time by a randomized algorithm) for the HITTING SET or the SET COVER problems associated with the corresponding range spaces.

Key words. geometric range spaces, ε -nets, Exponential Decay Lemma, set cover, hitting set

AMS subject classifications. 68W20, 68W25, 52C45, 68U05

DOI. 10.1137/090762968

1. Introduction. Since their introduction in 1987 by Haussler and Welzl [HW87] (see also Clarkson [Cla87] and Clarkson and Shor [CS89] for related techniques), ε -nets have become one of the central concepts in computational and combinatorial geometry and have been used in a variety of applications, such as range searching, geometric partitions, and bounds on curve-point incidences; see, e.g., Matoušek [Mat02]. We recall their definition: A *range space* (X, \mathcal{R}) is a pair consisting of an underlying universe X of objects and a certain collection $\mathcal{R} \subseteq 2^X$ of subsets (*ranges*). Of particular interest are range spaces of *finite VC-dimension*; the reader is referred to [HW87] for the exact definition. Informally, it suffices to require that, for any finite subset $P \subset X$, the number of distinct sets $r \cap P$ for $r \in \mathcal{R}$ be $O(|P|^d)$ for some constant d (which is upper-bounded by the VC-dimension of (X, \mathcal{R})).

Given a range space (X, \mathcal{R}) , a finite subset $P \subset X$, and a parameter $0 < \varepsilon < 1$, an ε -net for P and \mathcal{R} is a subset $N \subseteq P$ with the property that any range $r \in \mathcal{R}$ 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.

*Received by the editors June 24, 2009; accepted for publication (in revised form) April 13, 2010; published electronically July 29, 2010. A preliminary version of this paper appeared in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, ACM, New York, 2009, pp. 639–648.

<http://www.siam.org/journals/sicomp/39-7/76296.html>

[†]Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840 (aronov@poly.edu). This author’s work was supported by grant 2006/194 from the U.S.-Israel Binational Science Foundation, by NSA MSP grant H98230-06-1-0016, and by NSF grant CCF-08-30691.

[‡]Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (esther@courant.nyu.edu). Part of this author’s work was performed at the Department of Computer Science, Duke University, Durham, NC 27708-0129 and was supported by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376 and by NIH grant 1P50-GM-08183-01.

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (michas@post.tau.ac.il). This author’s work was supported by grant 2006/194 from the U.S.-Israel Binational Science Foundation, by NSF grants CCF-05-14079 and CCF-08-30272, by grants 155/05 and 338/09 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

The ε -net theorem of Haussler and Welzl asserts that, for any (X, \mathcal{R}) , P , and ε as above, such that (X, \mathcal{R}) has finite VC-dimension d , there exists an ε -net N of size $O\left(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon}\right)$, and that in fact a random sample of elements from P of that size is an ε -net with constant probability. In particular, the size of N is independent of the size of P . The bound on the size of the ε -net was later improved to $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ by Blumer et al. [BEHW89], and then to $(1 + o(1))\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}$ by Komlós, Pach, and Woeginger [KPW92].

In geometric applications, this abstract framework is used as follows. The ground set X is typically a set of simple geometric objects (points, lines, hyperplanes), and the ranges in \mathcal{R} are defined in terms of intersection with (or, for point objects, containment in) simply shaped regions (halfspaces, balls, simplices, etc.), formally assumed to be regions of *constant description complexity*, meaning that they are semialgebraic sets defined in terms of a constant number of polynomial equations and inequalities of constant maximum degree. It is known that in such cases the resulting range space (X, \mathcal{R}) does have finite VC-dimension (see, e.g., [SA95]).

For example, the main result of our paper concerns the range space in which the objects are points in the plane and the ranges are axis-parallel rectangles; more precisely, each range is the intersection of the ground set with such a rectangle. The *dual* range space in this case is one in which the objects are rectangles and each point p in the plane defines a range which is the subset of the given rectangles that contain p . An ε -net in this case is a subset of the rectangles that covers all the “deep” points.

One of the major questions in the theory of ε -nets, open since their introduction more than 20 years ago, is whether the factor $\log \frac{1}{\varepsilon}$ in the upper bound on their size is really necessary, especially in typical low-dimensional geometric situations. To be precise, in the general abstract context the answer is “yes,” as shown by Komlós, Pach, and Woeginger [KPW92], using a randomized construction on abstract hypergraphs (see also [PA95]). However, there is no known lower bound, better than the trivial $\Omega(1/\varepsilon)$, in any “concrete” case and certainly in any geometric situation of the kind mentioned above. The prevailing conjecture is that, at least in these geometric scenarios, there always exists an ε -net of size $O(1/\varepsilon)$ [MSW90].

This “linear” 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 pseudodisk ranges in the plane; see [MSW90, Mat92b, CV07, HKSS08, PR08]. Additional progress was made recently. Clarkson and Varadarajan [CV07], essentially adapting Matoušek’s technique [Mat92b] to their more general setting, have introduced a method for constructing small-size ε -nets in dual range spaces arising in geometric situations where, as above, the ground set is a collection of regions, and each point p determines a range equal to the set of those regions which contain p , and where the combinatorial complexity of the *union* of any finite number r of the regions in the ground set is small—specifically $o(r \log r)$. (The exact condition is slightly more involved; see below.) As a matter of fact, albeit not explicitly presented in this manner, the technique of [CV07] is more general and can also be applied to the primal version of the problem, provided that it satisfies a condition analogous to the one on small union complexity; see below for more details. More recently, Pyrga and Ray [PR08] have proposed a general abstract scheme for constructing small-size ε -nets in hypergraphs (i.e., range spaces) which satisfy certain properties and have applied it to the special cases of halfspaces in two and three dimensions and to several other related scenarios. Very recently, Varadarajan [Var09] has independently obtained a similar improvement on the bound of [CV07] for the size of an ε -net in the dual range

space of α -fat triangles and planar point sets using very different methods. Shortly after the publication of this work, King and Kirkpatrick [KK10] showed the existence of ε -nets of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$ for points in a simple polygon P , where the ranges are the visibility polygons from points on ∂P .

The SET COVER and HITTING SET problems. Given a range space (P, \mathcal{R}) , with P and \mathcal{R} finite, the SET COVER problem is to find a minimum-size subcollection $S \subseteq \mathcal{R}$, whose union covers P . A related (dual) problem is the HITTING SET problem, where we want to find a smallest-cardinality subset $H \subseteq P$, with the property that each range $r \in \mathcal{R}$ intersects H . Equivalently, a set cover for (P, \mathcal{R}) is a hitting set for the dual range space. The general (primal and dual) problems are NP-hard to solve (even approximately) [GJ79, Kar72], and the simple greedy algorithm yields the (asymptotically) best possible approximation factor of $O(1 + \log |P|)$ computable by a polynomial-time algorithm, under appropriate complexity-theoretic assumptions [BGLR93, Fei98].

Most of these problems remain NP-hard even in geometric settings [FG88, FPT81]. However, one can attain an improved approximation factor of $O(\log \text{OPT})$ in polynomial time for many of these scenarios, where OPT is the size of the optimal solution. This improvement is based on the techniques of Even, Rawitz, and Shahar [ERS05] and of Brönnimann and Goodrich [BG95] (see also Clarkson [Cla93]), where the key observation is the relation to ε -nets: The existence of an ε -net of size $O\left(\frac{1}{\varepsilon} \varphi\left(\frac{1}{\varepsilon}\right)\right)$ for any $\varepsilon > 0$ implies that these techniques generate, in expected polynomial time, a hitting set (or a set cover) whose size is $O(\text{OPT} \cdot \varphi(\text{OPT}))$.

Hence, for range spaces of finite VC-dimension, the Haussler–Welzl theorem leads to an approximation factor $O(\log \text{OPT})$. Consequently, improved bounds for the size of ε -nets, in the primal or the dual setting, imply improved approximation factors for the corresponding HITTING SET or SET COVER problems, at least in the context of randomized polynomial-time constructibility (which is what is provided by the procedures of Brönnimann and Goodrich or of Even, Rawitz, and Shahar).

Our results. In this paper we first consider the cases of point objects and axis-parallel rectangular ranges in the plane, and of point objects and axis-parallel box ranges in three dimensions, and show that both range spaces admit ε -nets of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$, thus significantly improving the standard bound $O\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$. Our technique is similar in spirit to those of Chazelle and Friedman [CF90] and of Clarkson and Varadarajan [CV07], but it differs from them in one key (and fairly simple) idea (that of *oversampling*; see below), which, incidentally, can also be used in the more general context of [CV07] to improve the bounds that are obtained there for the size of the respective ε -nets. We also propose a different probabilistic model, which eventually yields, as a by-product, a simpler analysis than that of [CV07]; see below. An interesting feature of our technique is that it can be extended to points and axis-parallel boxes in *any* dimension, provided that the input points are randomly and uniformly distributed in the unit cube.

We also describe how to construct these ε -nets in randomized expected nearly linear time. Our results then lead to randomized polynomial-time approximation algorithms for the HITTING SET problem in these two range spaces, involving axis-parallel rectangles and boxes, respectively, which guarantee an approximation factor of $O(\log \log \text{OPT})$.

In fact, we observe that the technique can be cast in a fairly general setting, which we present and analyze. This allows us to extend it, with relative ease, to many other concrete range spaces. Specifically, we extend it to the case of planar point sets and

TABLE 1

The bound on the ε -net size and the corresponding approximation factor for the HITTING SET problem for the primal range spaces of points and the regions listed in the table.

Dimension	Regions	ε -net size	Approximation factor
$d = 2$	axis-parallel rectangles	$O(\varepsilon^{-1} \log \log (\varepsilon^{-1}))$	$O(\log \log \text{OPT})$
$d = 2$	α -fat triangles	$O(\varepsilon^{-1} \log \log (\varepsilon^{-1}))$	$O(\log \log \text{OPT})$
$d = 3$	axis-parallel boxes	$O(\varepsilon^{-1} \log \log (\varepsilon^{-1}))$	$O(\log \log \text{OPT})$
$d \geq 1$	axis-parallel boxes with points uniformly distributed in $[0, 1]^d$	$O(\varepsilon^{-1} \log \log (\varepsilon^{-1}))$	$O(\log \log \text{OPT})$

TABLE 2

The bound on the ε -net size and the corresponding approximation factor for the SET COVER problem for the dual range spaces of the planar regions listed in the table and points.

Planar regions	ε -net size	Approximation factor
α -fat triangles	$O(\varepsilon^{-1} \log \log \log (\varepsilon^{-1}))$	$O(\log \log \log \text{OPT})$
locally γ -fat objects	$O(\varepsilon^{-1} \log \log (\varepsilon^{-1}))$	$O(\log \log \text{OPT})$
locally γ -fat objects of the same size	$O(\varepsilon^{-1} \log \beta_{s+2}(\varepsilon^{-1}))$	$O(\log \beta_{s+2}(\text{OPT}))$
semiunbounded pseudotrapezoids	$O(\varepsilon^{-1} \log \beta_{s+2}(\varepsilon^{-1}))$	$O(\log \beta_{s+2}(\text{OPT}))$
pseudo-halfplanes	$O(\varepsilon^{-1} \log \beta_s(\varepsilon^{-1}))$	$O(\log \beta_s(\text{OPT}))$
regions bounded by Jordan arcs with three intersections per pair	$O(\varepsilon^{-1} \log \alpha(\varepsilon^{-1}))$	$O(\log \alpha(\text{OPT}))$

α -fat triangles (for the primal range space), that is, triangles, each of whose angles is at least α , for some constant $\alpha > 0$ (see [MPSSW94]). In this case as well we show the existence of ε -nets of size $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$, leading to an approximation factor of $O(\log \log \text{OPT})$ for the corresponding HITTING SET problem. Table 1 summarizes these results.

Similarly, we obtain improved bounds for the size of ε -nets in the dual range space, and, consequently, for approximation factors for the corresponding SET COVER problem, in the following cases, all involving points and regions in the plane (refer to section 7 for the exact definition of the range spaces, and see Figure 14 for an illustration). We consider the following collections of planar regions: (a) α -fat triangles, (b) locally γ -fat objects (with slightly better bounds when all objects have nearly equal size), (c) semiunbounded pseudotrapezoids, and (d) regions bounded by Jordan arcs with three intersections per pair. These results are summarized in Table 2.

Our technique for rectangles—a brief overview. We start with a brief overview of our analysis, in which we assume some familiarity with the earlier papers [CF90, CV07] cited above. Let P be a given set of n points in the plane. We first sketch a somewhat simpler approach that *almost* works; it does not properly address a certain critical technical issue but captures the essence of our method. We then briefly describe how to modify it so that it does produce ε -nets of the desired size.

Put $r = 1/\varepsilon$. We draw a random sample R of $s \gg r$ points of P (this is the *oversampling* ingredient of our technique; the specific choice of s , made below, is crucial) and make R part of the ε -net to be constructed, and so it remains only to handle 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, so that any

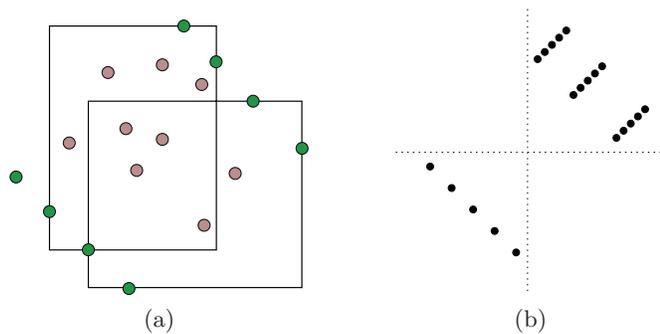


FIG. 1. (a) A configuration with quadratically many maximal R -empty rectangles (the points of R are shaded darker and lie on the two extreme staircases). (b) A configuration with an expected quadratic number of maximal R -empty rectangles, each containing $\Omega(n/s)$ points. The lower staircase contains $n/2$ points, and each of the s upper “diagonals” contains $\frac{n}{2s}$ points.

other R -empty rectangle is contained in one of them. 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, in a sense that we do not make very precise here, 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] and sample a small subset of $O(t \log t)$ points of $M \cap P$, where $t = s|M \cap P|/n$, to serve as a $(1/t)$ -net for $M \cap P$. The union of R and all these samples constitutes the desired ε -net; it is fairly easy to show that this is indeed an ε -net.

This approach does not quite work because, for a bad choice of R , the number of maximal R -empty rectangles can be $\Theta(s^2)$ in the worst case (see, e.g., [NLH84] and Figure 1(a)). Moreover, even if we consider only *random* subsets R , which is what the technique produces anyway, there are point sets where the *expected* number of maximal R -empty rectangles which contain $\Omega(n/s)$ points of P is still $\Theta(s^2)$; see Figure 1(b). Using the technique outlined above literally turns out to yield a bound of $\Theta\left(\frac{1}{\varepsilon^2}\right)$ on the expected size of the ε -net in the worst case, which is of course much too large.

We overcome this issue by modifying the scheme so that it produces fewer maximal empty rectangles. To do so, we reset $r := 2/\varepsilon$ and decompose the plane into a binary-tree-like hierarchy of vertical strips, stopping just before reaching strips that contain fewer than n/r points of P . For any rectangle \tilde{Q} which contains at least εn points of P , we find the first (highest in the hierarchy) strip-bounding line which crosses \tilde{Q} , take one of its halves, Q , which contains at least $\varepsilon n/2 = n/r$ points, and consider only such rectangles in the construction of our net. We thus face subproblems, each involving a vertical strip σ and the corresponding subset $P \cap \sigma$ of P , and ranges which are rectangles that are “anchored” at a specific side of σ (so that they effectively behave like three-sided unbounded rectangles for $P \cap \sigma$; refer to Figure 2). The number of maximal R -empty rectangles of this type, within σ , is only *linear* in $|R \cap \sigma|$, leading to an overall collection \mathcal{M} of maximal R -empty rectangles of the new kind, whose size is only $O(s \log r)$.

We now choose $s := cr \log r$ for an appropriate absolute constant c . Using

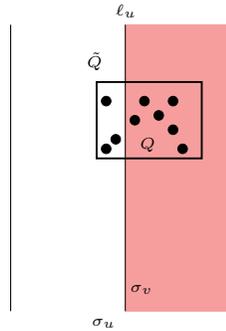


FIG. 2. The half-rectangle Q is anchored at the left entry side ℓ_u of the strip σ_v .

the so-called Exponential Decay Lemma of [AMS98, CF90], one can show that the expected number of maximal heavy empty rectangles that can contain rectangles Q of the above kind is only *sublinear* in r , which in turn implies that (even after taking into account the sizes of the resampled nets within each of these rectangles, and not just their number) the expected size of the ε -net is dominated by the expected size of R , namely, $O(r \log \log r) = O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$.

The general framework. In section 3 we argue that our technique can be applied in a more general setting. We formulate this extension in terms of an abstract framework in which the given ranges are required to satisfy several properties, similar to those which we used in the case of axis-parallel rectangles. When these properties hold, essentially the same machinery as in the case of rectangles can be applied to yield ε -nets of smaller size; see section 3 for complete details. We then apply this general technique to all the other primal range spaces listed above.

Improving the general bounds in [CV07]. Readers familiar with the technique of Clarkson and Varadarajan [CV07] will notice the similarity of our approach to theirs. The key new ingredient is that we use a *larger* initial sample R of expected size $\Theta(r \log \log r)$ rather than $O(r)$; that is, the essence of our approach is *oversampling*. The same idea can be applied in the more general context of [CV07] and leads to an improvement of each of their bounds that are superlinear in r . In a sense, the Clarkson–Varadarajan technique is a dual version of the general set-up discussed above for dual range spaces. Specifically, Clarkson and Varadarajan consider dual range spaces and show that if the union complexity of any m of the ranges (i.e., objects in the dual ground set) is $O(m\varphi(m))$ for an appropriate slowly increasing function φ , then there exist ε -nets in such a dual range space of size $O\left(\frac{1}{\varepsilon}\varphi\left(\frac{1}{\varepsilon}\right)\right)$. Using our approach, we obtain ε -nets of size $O\left(\frac{1}{\varepsilon} \log \varphi\left(\frac{1}{\varepsilon}\right)\right)$. Moreover, their method yields improved bounds for ε -nets only when $\varphi(m) = o(\log m)$, whereas our method yields improved bounds as long as $\varphi(m) = 2^{o(\log m)}$. The case of rectangles is interesting in this aspect because, with the addition of the divide-and-conquer decomposition scheme mentioned above, the complexity of the appropriate analogue of the union of m dual ranges (which is the number of maximal empty rectangles) is $O(m \log m)$, which is the threshold bound at which the more straightforward sampling approach of [CV07] fails.¹

¹As already noted, the $\log m$ factor comes from the binary-tree hierarchy—see what follows for details.

2. Small-size ε -nets for axis-parallel rectangles. Let P be a set of n points in the plane. Put $r := 2/\varepsilon$ and $s := cr \log \log r$, where $c > 1$ is an arbitrary constant. For simplicity of presentation, we assume that each of n , r is a power of 2, thereby avoiding the need for rounding. This involves no loss of generality. Construct a balanced binary tree \mathcal{T} over the points of P in their x -order in the following manner: Sort the points according to their x -coordinates; then group them into r groups, each of size n/r . Now construct a full binary tree of height $\log_2 r$ on these groups. The leaves of the tree correspond to the groups, and each internal node corresponds to the union of the groups stored at the leaves of its subtree. Each internal node stores a vertical line that separates the points of the left subtree from those of the right one. By construction, \mathcal{T} has $\log r$ levels.

Fix a random sample $R \subseteq P$, 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 s . The sample R is part of the ε -net N that we are about to construct.

Each node v of \mathcal{T} is associated with a subset P_v of P (resp., R_v of R), consisting of those points of P (resp., of R) stored at the subtree rooted at v . We also associate with v a vertical line ℓ_v which splits P_v into the two subsets P_{v_1}, P_{v_2} associated with the children v_1, v_2 of v . Using the lines ℓ_u , we associate with each node v a strip σ_v , which contains P_v (and R_v), where σ_{root} is the entire plane, and, for a left (resp., right) child node $v \neq \text{root}$ of its parent u , σ_v is the left (resp., right) portion of σ_u delimited by ℓ_u . We call ℓ_u the *entry side* of σ_v .

Note that, since the sets P_v are defined ahead of the draw of R , our sampling model guarantees that, for each node v , R_v is an unbiased sample of P_v , drawn from P_v by exactly the same rule, namely, by choosing each point independently with probability π .

Let \tilde{Q} be an axis-parallel rectangle containing at least εn points of P , and let u be the highest node of \mathcal{T} such that ℓ_u crosses \tilde{Q} , partitioning it into two parts, one of which necessarily contains at least $\varepsilon n/2 = n/r$ points of P . Denote that portion of \tilde{Q} by Q , and let v be the child of u such that $Q \subseteq \sigma_v$. We say that Q is *anchored* at the entry side ℓ_u of σ_v ; see Figure 2.

If Q contains a point of R , we are done, as $Q \subset \tilde{Q}$ and the goal was to construct a subset of P that meets every rectangle \tilde{Q} containing at least εn points of P . So we may assume that Q does not contain such a point; we then say that Q is *R -empty*; equivalently, Q is *R_v -empty*.

We define, for each node v of \mathcal{T} , a set \mathcal{M}_v consisting of all the maximal (open) anchored R_v -empty axis-parallel rectangles contained in σ_v . Without loss of generality, assume that the entry side ℓ_u of σ_v is its left side. In general, a rectangle M in \mathcal{M}_v is determined by three points of R_v , one point lying on each of the three unanchored sides of M (see Figure 3(a)), but \mathcal{M}_v may also contain degenerate rectangles M where some (or all) of these points are missing, in which case M extends as much as possible, within σ_v , in the appropriate direction (upward, downward, or to the right). In particular, when $R_v = \emptyset$, there is precisely one maximal R_v -empty rectangle, namely, the whole strip; see Figure 3(b)–(e), which illustrates some of these cases.

It is easy to show that $|\mathcal{M}_v| = 2r_v + 1$, where $r_v := |R_v|$. Indeed, if a rectangle M has a point $q \in R_v$ on its right side, then q cannot lie on the right side of any other rectangle in \mathcal{M}_v , so the number of such rectangles is r_v (equality is also easy to verify). Otherwise, the points of R_v on the top and bottom sides of M must be consecutive in R_v in the y -order, and there are $r_v - 1$ such pairs. Finally, there are two semiunbounded rectangles—one delimited from below by the highest point of R_v

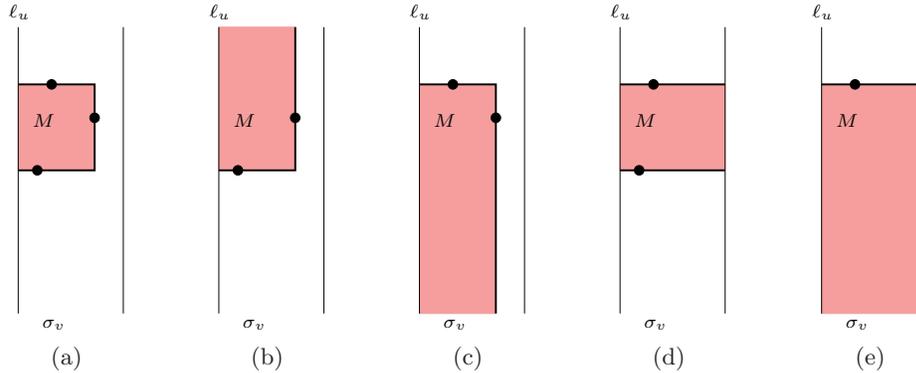


FIG. 3. An anchored maximal R -empty rectangle that is determined by three points (a), by a pair of points (b)–(d), or by a single point (e).

and the other delimited from above by the lowest point (as in Figure 3(e)). It is easily checked that the bound $2r_v + 1$ also applies when $r_v = 0, 1$. It thus follows that the overall number of such maximal empty rectangles $M \in \mathcal{M}_v$, over all nodes v of \mathcal{T} at any fixed level, is $O(|R| + r')$, where r' is the number of nodes at the level, and the total over all levels of \mathcal{T} is $O(r + |R| \log r)$.

Returning now to the anchored rectangle Q and the corresponding node v , we note that Q is contained in at least one rectangle in \mathcal{M}_v . Indeed, assuming, as above, that the entry side of σ_v is its left side, expand Q by pushing its right side to the right until it touches a point of R_v or reaches the right side of σ_v , and then push the top and bottom sides until each of them meets a point of R_v or extends to $\pm\infty$. The resulting rectangle belongs to \mathcal{M}_v and encloses Q .

For each node v of \mathcal{T} , and each member $M \in \mathcal{M}_v$, define the *weight factor* t_M of M to be $\lceil s|M \cap P|/n \rceil$. 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 anchored rectangle Q , as above, can be completely contained in one of them. By the standard ε -net theory [HW87], for each $M \in \mathcal{M}_v$ with $t_M \geq c \log \log r$, there exists a subset $N_M \subseteq M \cap P_v$ of size $c't_M \log t_M$ that forms a $(1/t_M)$ -net for $M \cap P_v$, where c' is another absolute constant.

The final ε -net N is the union of R with the sets N_M , over all the heavy rectangles M (i.e., rectangles with $t_M \geq c \log \log r$) in the respective sets \mathcal{M}_v , over all nodes v of \mathcal{T} .

N is an ε -net. Since $R \subseteq N$, it suffices to show that for any R -empty rectangle Q contained in a strip σ_v , anchored at the entry side of σ_v , and containing at least $\varepsilon n/2 = n/r$ points of P (i.e., of P_v), and for any $M \in \mathcal{M}_v$ which contains Q , we have $Q \cap N_M \neq \emptyset$. We have

$$\frac{|Q \cap P|}{|M \cap P|} \geq \frac{n/r}{nt_M/s} = \frac{c \log \log r}{t_M} \geq \frac{1}{t_M}.$$

Since N_M is a $(1/t_M)$ -net for $M \cap P$, it follows that $Q \cap N_M \neq \emptyset$, as asserted. Note that the above inequality implies that we need not sample that many points in N_M and can make do with $c't_M^* \log t_M^*$ points, where $t_M^* := t_M/(c \log \log r)$. However, this slight improvement does not asymptotically affect the bound that we are about to derive.

Estimating the size of N . The expected size of N is equal to

$$\begin{aligned} & \mathbf{Exp} \left\{ |R| + c' \sum_v \sum_{\substack{M \in \mathcal{M}_v \\ t_M \geq c \log \log r}} t_M \log t_M \right\} \\ &= cr \log \log r + c' \cdot \mathbf{Exp} \left\{ \sum_v \sum_{\substack{M \in \mathcal{M}_v \\ t_M \geq c \log \log r}} t_M \log t_M \right\}. \end{aligned}$$

We continue the analysis using the notation of [AMS98]. Fix a level i ; each node v at this level satisfies $|P_v| = n/2^i$. Let $\text{CT}(R)$ denote the union of the collections \mathcal{M}_v over all nodes v at level i . For a positive parameter t , let $\text{CT}_t(R)$ denote the subset of $\text{CT}(R)$ consisting of those rectangles M with $t_M \geq t$. Let R' denote another random sample of P , where each point $p \in P$ is now chosen independently to belong to R' with probability $\pi' := \pi/t$.

Let \mathcal{C} denote the set of all rectangles M such that M is anchored at the entry side of σ_v for some node v at level i and has one point of P on each of its three other sides (the cases of degenerate rectangles, determined by fewer than three points, are treated in a fully analogous manner). For a rectangle $M \in \mathcal{C}$, its *defining set* $D(M)$ is the set of these three points, and its *killing set* $K(M)$ is the set of points of P in the interior of M . (Recall that throughout this discussion we have fixed the level i .)

Agarwal, Matoušek, and Schwarzkopf [AMS98] impose two axioms on the sets $\text{CT}(R)$. These axioms are too intricate for what we need here, while they are necessary to handle the more involved scenario considered in [AMS98]. For our purpose, we can replace them by the single “axiom” asserting that a rectangle $M \in \mathcal{C}$ belongs to $\text{CT}(R)$ if and only if $D(M) \subseteq R$ and $K(M) \cap R = \emptyset$, which holds by construction in our setting. (We also caution the reader that our sampling model is different from that of [AMS98]; they pick a random subset of a fixed given size uniformly from all such subsets, whereas we independently choose each point of P to belong to the sample. Nevertheless, the lemma, given below, also holds in our model; if anything, the analysis is simpler. For the sake of completeness, we give, in the appendix, a short (but complete) proof of our variant of the lemma.)

LEMMA 2.1 (Exponential Decay Lemma; Agarwal, Matoušek, and Schwarzkopf [AMS98]).

$$\mathbf{Exp} \{ |\text{CT}_t(R)| \} = O(2^{-t} \mathbf{Exp} \{ |\text{CT}(R')| \}).$$

We apply the lemma with $t = c \log \log r$, so $\pi' = \pi/t = r/n$. Recall that $\text{CT}(R')$ is the set of all maximal R' -empty rectangles anchored at the entry sides of their respective strips σ_v at the fixed level i . Their number is $|\text{CT}(R')| = \sum_v (2r'_v + 1)$, where $R'_v := R' \cap \sigma_v$ and $r'_v := |R'_v|$. Since the sets R'_v at level i are disjoint, $\sum_v r'_v = |R'|$. Hence, since there are at most $2r$ nodes at a fixed level of the tree, we have $|\text{CT}(R')| \leq 2|R'| + 2r$. Hence $\mathbf{Exp} \{ |\text{CT}(R')| \} = O(r)$. We thus have

$$\mathbf{Exp} \{ |\text{CT}_t(R)| \} = O(2^{-t} \mathbf{Exp} \{ |\text{CT}(R')| \}) = O(r2^{-c \log \log r}) = O(r/\log^c r).$$

More generally, for any $j \geq t$, we have $\mathbf{Exp} \{ |\text{CT}_j(R)| \} = O(r/2^j)$, as is easily checked.

Getting back to the contribution of the fixed level i to the expected size of N , we

have (where $t = c \log \log r$)

$$\begin{aligned}
 \mathbf{Exp} \left\{ \sum_{v \text{ at level } i} \sum_{\substack{M \in \mathcal{M}_v \\ t_M \geq t}} t_M \log t_M \right\} &= \mathbf{Exp} \left\{ \sum_{j \geq t} \sum_{\substack{M \in \text{CT}(R) \\ t_M = j}} j \log j \right\} & (*) \\
 &= \mathbf{Exp} \left\{ \sum_{j \geq t} j \log j \cdot (|\text{CT}_j(R)| - |\text{CT}_{j+1}(R)|) \right\} \\
 &= \mathbf{Exp} \left\{ t \log t \cdot |\text{CT}_t(R)| \right. \\
 &\quad \left. + \sum_{j > t} (j \log j - (j-1) \log(j-1)) |\text{CT}_j(R)| \right\} \\
 &= O \left(\frac{r}{\log^c r} (t \log t) + \sum_{j > t} \frac{r}{2^j} \log j \right) \\
 &= O \left(\frac{rt \log t}{\log^c r} \right) = O \left(\frac{r \log \log r \log \log \log r}{\log^c r} \right).
 \end{aligned}$$

Recall again that the analysis so far has been confined to a single level i . Repeating it for each of the $1 + \log r$ levels, we obtain, since $c > 1$,

$$\mathbf{Exp} \{|N|\} = O \left(r \log \log r + \frac{r \log \log r \log \log \log r}{\log^{c-1} r} \right) = O(r \log \log r).$$

We have thus shown the following theorem.

THEOREM 2.2. *For any set P of n points in the plane and a parameter $\varepsilon > 0$, there exists an ε -net of P , of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$, for axis-parallel rectangles.*

Remark. A key ingredient of the analysis is that we have managed to reduce the expected number of R -empty rectangles from $\Theta(s^2)$ to $O(s \log r)$, using a decomposition of the point set into canonical subsets, so that (i) any rectangle Q with at least εn points of P interacts with just *two* subsets (any constant number would do just as well), and (ii) for each canonical subset, the number of maximal R -empty rectangles (now anchored at the entry side of the respective strip and fully contained in that strip) is only linear in the number of sample points in that strip.

Constructing the ε -net. We next present a randomized algorithm for constructing an ε -net of the above size.

We construct the balanced binary tree \mathcal{T} over the points of P in $O(n \log r)$ time (stopping at nodes v for which $|P_v| = n/r$) and generate the random sample R using the drawing model assumed above; the expected size of R is s .

Following the above notation, we associate with each node $v \neq \text{root}$ of \mathcal{T} a strip σ_v , the subsets P_v, R_v , and an entry side ℓ_u of σ_v (where u is the parent of v). We next construct, for each such node v , the set \mathcal{M}_v of all maximal anchored R_v -empty axis-parallel rectangles contained in σ_v . This is easy to do in time $O(r_v \log r_v)$, where $r_v := |R_v|$, as follows. Assume that ℓ_u is the left side of σ_v . Sort the points of R_v by their y -coordinates, and find, for each point q , the lowest point q' which lies above q and to its left. This can be done in linear time by scanning the points of R_v in decreasing y -order and by dynamically maintaining the sorted sequence of xy -minima [CLRS01]. Symmetrically, we find, for each point q , the highest point q'' which lies below q and to its left. The resulting triples (q, q', q'') (including degenerate ones) determine r_v of the maximal empty rectangles in \mathcal{M}_v . Each of the other $r_v + 1$

rectangles straddles σ_v from left to right and either is delimited by a pair of points of R_v , consecutive in the y -order, which lie on its top and bottom sides, or is an unbounded half-strip, bounded by a single point.

It thus follows that the overall expected running time for constructing the sets \mathcal{M}_v , over all nodes v at a fixed level i , is $O(s \log r)$ (by bounding $\log r_v$ by $\log r$ and using linearity of expectations) for a total of $O(s \log^2 r)$ time over all levels i .

We next count, for each resulting rectangle M , the number of points in $M \cap P$ using a standard two-dimensional range-tree data structure. This yields the respective weight factors t_M , as defined above; we keep only those rectangles with $t_M \geq c \log \log r$. For each of these surviving rectangles M , we *report* the set $P \cap M$ and construct a $(1/t_M)$ -net for $P \cap M$ using, e.g., the deterministic algorithm of Matoušek [Mat95] (or a straightforward random sampling mechanism [HW87]). We output the union of R with all the resulting nets. Using the Exponential Decay Lemma and considerations similar to those in the proof of Theorem 2.2, it can be shown that the overall expected number of reported points in the sets $P \cap M$, over all heavy rectangles M and nodes v , is only linear in n .

As argued above, the output N is guaranteed to be an ε -net for P (if we construct the subnets N_M deterministically). The size of N is a random variable whose expectation is $O(r \log \log r)$. We can ensure this size with high probability by discarding outputs that are too large and by repeating the sampling.

The entire algorithm takes $O(n \log n)$ randomized expected time, as is easily seen.

Remark. The running time of the algorithm can be slightly improved to $O(n \log r)$ using machinery similar to that used above, where we modify the range-tree data structure to support only *approximated* range counting queries (where the approximation is within a factor of roughly n/r). In this case the height of the range-tree decreases to $O(\log r)$. We leave the easy details to the reader.

3. An abstract framework. In this section we generalize the method used above in order to apply it in several other situations.

We are given a range space (P, \mathcal{R}) , where P is a finite set of points in \mathbf{R}^d and where \mathcal{R} is some class of d -dimensional objects of constant description complexity. We seek conditions on \mathcal{R} that would enable us to extend the technique of the previous section to construct small ε -nets for (P, \mathcal{R}) .

We call a region in \mathbf{R}^d Y -empty if its interior does not intersect Y for a set $Y \subseteq \mathbf{R}^d$.

DEFINITION 3.1. *We say that (P, \mathcal{R}) has the small cover property if there exist two other classes, $\mathcal{R}_0, \mathcal{R}_1$, each consisting of d -dimensional objects of constant description complexity, so that each object of \mathcal{R}_1 is defined by a constant number of points of P in the sense of [CV07, CS89]. We also assume that there exist constants $A, B \geq 1$, so that the following conditions hold. Consider any $\varepsilon > 0$, $1 \leq s \leq n$, and any random sample $R \subseteq P$, obtained by picking each element of P independently with probability $\pi := s/n$. Let \mathcal{R}'_0 be the subcollection of \mathcal{R}_0 consisting of the R -empty regions in \mathcal{R}_0 , and let \mathcal{R}'_1 be the subcollection of \mathcal{R}_1 consisting of the R -empty regions in \mathcal{R}_1 , such that all their defining points belong to R . Then the following hold.*

(a) *For every R -empty region $Q \in \mathcal{R}$, with $|Q \cap P| > B\varepsilon n$, there exists $Q_0 \in \mathcal{R}'_0$ such that*

- (i) $Q_0 \cap P \subset Q \cap P$,
- (ii) $|Q_0 \cap P| > \varepsilon n$, and
- (iii) Q_0 is covered by at most A elements of \mathcal{R}'_1 , or, more precisely, $Q_0 \cap P \subseteq \bigcup_i (Q_1^{(i)} \cap P)$ for some choice of at most A sets $Q_1^{(i)} \in \mathcal{R}'_1$.

(b) *The expected size of \mathcal{R}'_1 , over the random choices of R , is at most $s\varphi(s)$, where $\varphi(x)$ is a sublinear function (i.e., one satisfying $\varphi(xy) \leq x\varphi(y)$ for any integers $x, y \geq 1$) and s is the expected size of R .*

Note that this set-up does indeed extend that for points and axis-parallel rectangles. Here \mathcal{R}_0 is the set of all axis-parallel anchored rectangles over all strips in the binary decomposition, \mathcal{R}_1 is the set of anchored axis-parallel rectangles, each defined by up to three points on its boundary, in the manner described in section 2, again, over all strips, $A = 1$, $B = 2$, and $\varphi(x) = O(\log x)$.

THEOREM 3.2. *If (P, \mathcal{R}) has the small cover property, with the parameters as specified in Definition 3.1, then it admits an ε -net of size $O(\frac{1}{\varepsilon} \log \varphi(\frac{1}{\varepsilon}))$ for any $0 < \varepsilon \leq 1$.*

Proof. We follow the same general construction as for the case of points and axis-parallel rectangles in the plane. Put $r := B/\varepsilon$, $s := Cr \log \varphi(r)$, and $\pi := s/n$, where $C > A$ is any fixed constant. We draw a random sample R of points of P , picking each point, independently, with probability π .

We then form the sets $\mathcal{R}'_0 \subseteq \mathcal{R}_0$ and $\mathcal{R}'_1 \subseteq \mathcal{R}_1$ of R -empty regions; by assumption, they satisfy the above properties. We define the weight factor t_M of a region $M \in \mathcal{R}'_1$ to be $s|M \cap P|/n$. By the standard ε -net theory [HW87], it follows that, with high probability, we have $|M \cap P| = O(\frac{n}{s} \log s)$ for each region M , and, in an informal and imprecise sense, the expected size of $M \cap P$, for a region M , is only $O(n/s)$. As above, we take each “heavy” region $M \in \mathcal{R}'_1$, with $|M \cap P| \geq n/(Ar)$, or, equivalently, with $t_M \geq s/(Ar) = (C/A) \log \varphi(r)$, and use the standard theory of ε -nets to deduce that there exists a $(1/t_M)$ -net N_M for the range space $(M \cap P, \mathcal{R}'_0)$, whose size is $O(t_M \log t_M)$. We output the union of R with all the sets N_M , over all heavy regions M , as the desired ε -net N .

Adapting the arguments from section 2, it is easy to verify that N is an ε -net for (P, \mathcal{R}) . Indeed, let Q be a range of \mathcal{R} containing at least $\varepsilon n = Bn/r$ points of P and not containing any point of R . Then Q contains a subrange $Q_0 \in \mathcal{R}'_0$ which contains at least n/r points of P . By assumption, there exist at most A sets of \mathcal{R}'_1 which cover $Q_0 \cap P$, so at least one of them, M , satisfies $|M \cap Q_0 \cap P| \geq n/(Ar)$. By the choice of t_M , we have

$$\frac{|M \cap P \cap Q_0|}{|M \cap P|} \geq \frac{n/(Ar)}{nt_M/s} = \frac{1}{t_M} \cdot \frac{s}{Ar} \geq \frac{1}{t_M},$$

so $M \cap Q_0$ must contain a point of N_M , showing that N is indeed an ε -net, as claimed.

To bound the expected size of N , we follow the previous analysis and apply the Exponential Decay Lemma in the new set-up, where $\text{CT}(R)$ is our set \mathcal{R}'_1 , and $\text{CT}_t(R)$ is the subset of $\text{CT}(R)$ consisting of those regions with weight factor at least t ; the defining set and the killing set are defined in a manner similar to that used earlier.

It thus follows that the Exponential Decay Lemma is applicable in this scenario as well, and it implies that, for any t ,

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

where R (resp., R') is a random sample in which each point of P is chosen independently with probability s/n (resp., $s/(tn)$).

To bound the latter expectation, we argue as follows.² Let $z := s/t$ denote the expected value of $|R'|$. By Chernoff's bound (see, e.g., [AS92]),

$$\Pr \{|R'| \geq \xi z\} \leq e^{-(\xi-1)^2 z/3}$$

²Here we pay back a little for using the simpler sampling model.

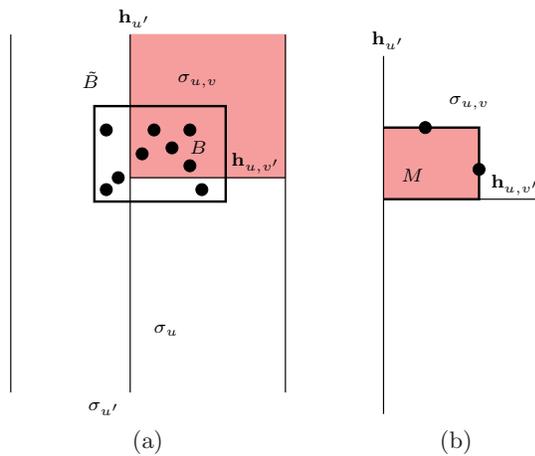


FIG. 4. A two-dimensional illustration: (a) The box B is anchored at the (apex of the) quadrant $\sigma_{u,v}$ (octant in 3-space). (b) An anchored box that is determined by a pair of points (a triple in 3-space).

for any $\xi > 1$. Hence, using the sublinearity of φ ,

$$\begin{aligned} \mathbf{Exp} \{ |R'| \varphi(|R'|) \} &\leq 2z\varphi(2z) + \sum_{j \geq 2} \Pr \{ |R'| \geq jz \} (j+1)z\varphi((j+1)z) \\ &\leq z\varphi(z) \cdot \left(4 + \sum_{j \geq 2} (j+1)^2 e^{-(j-1)^2 z/3} \right) = O(z\varphi(z)). \end{aligned}$$

In particular, for $t = c \log \varphi(r)$, each point is chosen in R' with probability r/n (so $z = r$), and we get

$$\mathbf{Exp} \{ |CT_t(R)| \} = O \left(2^{-c \log \varphi(r)} r \varphi(r) \right) = O \left(\frac{r}{\varphi^{c-1}(r)} \right),$$

which, for $c > 1$, is sublinear in r . The expectation is $O(2^{-t}(s/t)\varphi(s/t))$ for larger values of t .

We can now continue with the original analysis almost verbatim, arguing that the overall expected size of the subsamples “within” each heavy region of \mathcal{R}'_1 is sublinear in r , so the expected size of N is dominated by that of R ; thus it is $O(r \log \varphi(r))$. \square

In the following sections, we present several cases in which Theorem 3.2 can be applied, leading to an improved bound on the size of the corresponding ε -nets. With a few exceptions, the heart of the analysis is in finding good sets \mathcal{R}_0 and \mathcal{R}_1 which satisfy the conditions in the above definition.

4. Small-size ε -nets for axis-parallel boxes in three dimensions. We next extend our construction to the three-dimensional case, drawing upon the general construction provided in Theorem 3.2. We now let P be a set of n points in \mathbf{R}^3 and \mathcal{R} be the set of all axis-parallel boxes, and we put $r := 8/\varepsilon$ and $s := cr \log \log r$ for some fixed constant $c > 3$. We use a sampling model similar to that used in the two-dimensional problem in order to generate a random subset $R \subseteq P$ of expected size s . As above, we simplify the presentation by assuming n and r to be powers of 2.

We now describe the construction of the sets \mathcal{R}_0 and \mathcal{R}_1 and the various other parameters in the preceding analysis. To this end, we construct a three-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, by their y -coordinates in each secondary tree, and by their z -coordinates in each tertiary tree. We associate with each node u of the primary tree the subset P_u of points that it represents and a secondary (y -sorted) tree \mathcal{T}_u on P_u . Similarly, with each node v of a secondary tree \mathcal{T}_u we associate the corresponding subset $P_{u,v}$ of P_u and a tertiary (z -sorted) tree $\mathcal{T}_{u,v}$. Finally, each node w of a tertiary tree $\mathcal{T}_{u,v}$ is associated with the corresponding subset $P_{u,v,w}$ of $P_{u,v}$. We construct each of the three levels of \mathcal{T} down to nodes for which the size of their associated subset is n/r . Clearly, each of the primary, secondary, and tertiary trees has at most $\log r$ levels, and the total number of nodes in the range-tree \mathcal{T} is $O(r \log^2 r)$. Moreover, the sum of the sizes of all the subsets stored at the various nodes is $O(n \log^3 r)$; see, e.g., [dBCKO08] for further (standard) details.

Similar to the set-up in section 2, we associate with each nonleaf node of any subtree an axis-parallel plane which evenly splits the subset stored at the node into the two subsets stored at its children. More specifically, each nonleaf node u of the primary tree stores a plane \mathbf{h}_u orthogonal to the x -axis, each nonleaf node v of a secondary tree \mathcal{T}_u stores a plane $\mathbf{h}_{u,v}$ orthogonal to the y -axis, and each nonleaf node w of a tertiary tree $\mathcal{T}_{u,v}$ stores a plane $\mathbf{h}_{u,v,w}$ orthogonal to the z -axis.

These planes define, for each node w of a tertiary tree $\mathcal{T}_{u,v}$, an octant $\sigma_{u,v,w}$ which is the intersection of three halfspaces $H_u \cap H_{u,v} \cap H_{u,v,w}$, where (i) H_u is the halfspace bounded by $\mathbf{h}_{u'}$ and containing P_u , where u' is the parent of u ; (ii) $H_{u,v}$ is the halfspace bounded by $\mathbf{h}_{u,v'}$ and containing $P_{u,v}$, where v' is the parent of v in \mathcal{T}_u ; and (iii) $H_{u,v,w}$ is the halfspace bounded by $\mathbf{h}_{u,v,w'}$ and containing $P_{u,v,w}$, where w' is the parent of w in $\mathcal{T}_{u,v}$. In what follows we consider only triples (u, v, w) of vertices, each of which has a parent in its respective tree. Thus all three halfspaces are proper, and $\sigma_{u,v,w}$ is a nondegenerate octant. (Note, though, that, in general, it is more accurate to regard $\sigma_{u,v,w}$ as a box, or a clipped octant, bounded on the other sides also by planes associated with ancestors of u, v , and w . Nevertheless, in most of the following analysis, it suffices to treat $\sigma_{u,v,w}$ as an octant.)

Let B_0 be an axis-parallel box containing at least $\varepsilon n = 8n/r$ points of P . Let u' be the highest node in \mathcal{T} , so that the plane $\mathbf{h}_{u'}$ meets B_0 . This plane partitions B_0 into two portions, one of which (call it B_1) contains at least $4n/r$ points of P . Let u be the corresponding child of u' so that H_u contains B_1 . Next, let v' be the highest node in \mathcal{T}_u , such that $\mathbf{h}_{u,v'}$ meets B_1 , partitioning it into two portions, one of which, B_2 , contains at least $2n/r$ points of P . Let v be the child of v' for which $H_{u,v'}$ contains B_2 . Finally, let w' be the highest node in $\mathcal{T}_{u,v}$, such that $\mathbf{h}_{u,v,w'}$ meets B_2 , partitioning it into two portions, one of which, B , contains at least n/r points of P . Let w be the child of w' for which $H_{u,v,w'}$ contains B . (Note that u, v , and w are well defined in the sense that each of the subboxes is indeed split by a plane associated with a node in the corresponding truncated tree and does not reach a leaf without being split.)

By construction, B is *anchored* at the resulting octant $\sigma := \sigma_{u,v,w}$ in the sense that the apex o of σ is a vertex of B and the three facets of B adjacent to o lie on the three respective axis-parallel planar quadrants bounding σ . See Figure 4(a) for an illustration of (the two-dimensional analogue of) this scenario. We take \mathcal{R}_0 to be the set of all R -empty anchored boxes of this kind over all triples u, v, w .

For each node w of a tertiary tree $\mathcal{T}_{u,v}$, put $R_{u,v,w} = R \cap \bar{\sigma}_{u,v,w}$, where $\bar{\sigma}_{u,v,w}$ is the actual box that the ‘‘octant’’ $\sigma_{u,v,w}$ represents (see the comment above), and $r_{u,v,w} = |R_{u,v,w}|$. Let $\mathcal{M}_{u,v,w}$ denote the set of all maximal anchored R -empty (i.e., $R_{u,v,w}$ -empty) axis-parallel boxes contained in the octant $\sigma_{u,v,w}$. Since each box

$M \in \mathcal{M}_{u,v,w}$ behaves as an octant inside $\sigma_{u,v,w}$, it is determined by at most three points of $R_{u,v,w}$, each lying on a distinct facet of M ; see Figure 4(b) for a two-dimensional illustration. The number of such empty boxes (or, rather, octants) is only $O(r_{u,v,w} + 1)$, as shown³ in [BSTY98, KRSV08]. It thus follows that the size of the combined set $\mathcal{M} = \bigcup_{u,v,w} \mathcal{M}_{u,v,w}$, over all nodes w of all tertiary trees $\mathcal{T}_{u,v}$, is $O(|R| \log^3 r + r \log^2 r)$.

We are now in a position to apply Theorem 3.2 to the set-up just developed, with the collections \mathcal{R}_0 as just defined and $\mathcal{R}_1 = \bigcup_{u,v,w} \mathcal{R}_{u,v,w}$, where $\mathcal{R}_{u,v,w}$ is the set of all anchored axis-parallel boxes contained in the octant $\sigma_{u,v,w}$ (and then \mathcal{R}'_1 is set to \mathcal{M}). Clearly, all the required properties hold. In particular, every $(8n/r)$ -heavy and R -empty range of \mathcal{R} contains an (n/r) -heavy range of \mathcal{R}'_0 , and every range of \mathcal{R}'_0 is contained in a region of $\mathcal{M} = \mathcal{R}'_1$. Here $A = 1$ and $B = 8$. The size of \mathcal{M} , as noted above, is $O(|R| \log^3 r + r \log^2 r)$, whose expected value is at most $O(s \log^3 r)$. We thus have the following theorem.

THEOREM 4.1. *For any set P of n points in \mathbf{R}^3 and a parameter $\varepsilon > 0$, there exists an ε -net of P , for axis-parallel boxes, of size $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$.*

Constructing the ε -net. We construct an ε -net of this size using an easy extension of the algorithm presented in section 2. We start by building a three-level range-tree over the points of P , using $O(n \log^2 n)$ time and storage. The enumeration of the maximal anchored $R_{u,v,w}$ -empty octants in any canonical octant $\sigma_{u,v,w}$ can be performed in $O(|R_{u,v,w}| \log^2 r)$ time, using the algorithm described in [KRSV08]. Using our range-tree, we compute the weight factor t_M of each maximal octant M , collect all the heavy octants M (using counting queries), report the corresponding subsets $P \cap M$, and construct, for each such octant M , a $(1/t_M)$ -net of size $O(t_M \log t_M)$ for $P \cap M$ using standard techniques, as in the two-dimensional case. Omitting the further easy details, we obtain that the expected running time of the algorithm is $O(n \log^2 n)$, as asserted.

As in the planar case, the algorithm can be slightly improved to $O(n \log^2 r)$ using similar refinements.

5. Random point sets in any dimension. The technique fails in four and higher dimensions because the number of maximal empty orthants with respect to a set of m points can be $\Theta(m^{\lfloor d/2 \rfloor})$ (see [BSTY98, KRSV08]), which is at least quadratic for $d \geq 4$. It is a challenging open problem to extend our results to points and axis-parallel boxes in four and higher dimensions.

Nevertheless, there is one scenario in which the technique works in any dimension, which is the case when the ground set P consists of randomly and uniformly distributed points in \mathbf{R}^d . Specifically, we assume that each point of P is chosen independently at random from the uniform distribution in $[0, 1]^d$. As shown in [KRSV08], the expected number of maximal empty boxes in this case, for a set of m points, is only $O(m \log^{d-1} m)$ (see also [NLH84] for an earlier treatment of the planar case). Moreover, our random sampling model (where the random choices are assumed, of course, to be made independently of the random drawings of the points of the input set) ensures that the sample R is also an unbiased set of randomly, independently, and uniformly distributed points, so the expected number of maximal R -empty boxes is $O(s \log^{d-1} s)$ (this is proved as a special case of the analysis in section 3); the ex-

³In fact, the result in [KRSV08] is more general. It asserts that the number of maximal empty orthants for a set of m points in \mathbf{R}^d is $O(m^{\lfloor d/2 \rfloor})$. It is the nonlinearity of this bound for $d \geq 4$ which hampers the extension of our technique to higher dimensions.

pectation is with respect to both the random drawing of the points of the input set and our drawing of the sample R .

The set-up that we have reached is similar to those obtained in two and three dimensions. Thus Theorem 3.2 can be applied, and we obtain the following theorem.

THEOREM 5.1. *For any set P of n points in \mathbf{R}^d , each of which is drawn independently from the uniform distribution on $[0, 1]^d$, and a parameter $\varepsilon > 0$, there exists an ε -net of P , for axis-parallel boxes, of expected size⁴ $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$.*

6. Small-size ε -nets for fat triangles in the plane. In this section we present an extension of our technique to the range space of points in the plane and α -fat triangles for some fixed constant $\alpha > 0$. Recall that a triangle is α -fat if each of its angles is at least α . We thus have a set P of n points in the plane, and a parameter $\varepsilon > 0$, and our goal is to construct a small-size ε -net $N \subseteq P$ so that any α -fat triangle that contains at least εn points of P contains a point of N . As before, we assume that n , and the parameter r introduced below, are powers of 2.

Passing to semicanonical triangles. Following the analysis of [MPSSW94], we cover each α -fat triangle T by a triple of “semicanonical” $(\alpha/2)$ -fat triangles, each of which has a pair of edges with orientations taken from a fixed finite set \mathcal{D} of $O(1/\alpha)$ directions, and a third edge that bounds T ; see [MPSSW94, Lemma 3.2] and Figure 5(a). Clearly, if T contains at least εn points of P , then at least one of the three covering triangles contains at least $\varepsilon n/3$ points of P .

This canonization step yields a constant number ($O(1/\alpha^3)$, to be precise) of subfamilies of $(\alpha/2)$ -fat triangles, where the triangles in each subfamily have two edges at fixed orientations in \mathcal{D} and a third edge whose orientation can be assumed to belong to a sufficiently small range. Our strategy is thus to construct an $(\varepsilon/3)$ -net for P and each of these subfamilies, and the union of all these nets will be an ε -net for P and the family of all α -fat triangles.

Thus, in what follows we focus on a fixed semicanonical family \mathcal{F} . As in [MPSSW94], by applying an appropriate affine transformation, we may assume that each triangle $T \in \mathcal{F}$ is an “almost isosceles” right triangle with one horizontal edge and one vertical edge, which meet at the lower-left vertex of T , so that the orientation of the hypotenuse of T differs from 135° by at most one degree, say; see Figure 5(b).

Thus let P and \mathcal{F} be as above, and put $r := 132/\varepsilon$ and $s := cr \log \log r$ for some fixed constant $c > 2$.

We begin the construction of our ε -net by preparing an auxiliary $(1/r)$ -net N_1 of size $O(r \log \log r) = O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$ for axis-parallel rectangular ranges, as provided by Theorem 2.2, and include it in our overall net. Thus, if a triangle $T \in \mathcal{F}$ contains a “heavy” rectangle, which contains at least n/r points of P , then it also contains a point of N_1 , and we are done. We may therefore restrict the analysis to triangles $T \in \mathcal{F}$ which do not contain any heavy axis-parallel rectangle.

We use a sampling model similar to that used in the cases of axis-parallel rectangles and boxes for drawing a random subset $R \subseteq P$ of expected size s . We include R in our ε -net, so we need only cater to R -empty triangles $T \in \mathcal{F}$.

We next construct a two-level range-tree \mathcal{T} , over the points of P , in a manner analogous to that presented in section 4. The points are sorted by their x -coordinates in the primary tree and by their y -coordinates in each secondary tree, and we construct each of the two levels of \mathcal{T} down to nodes for which the size of their associated subset is n/r . (We will shortly add a third level to \mathcal{T} to handle one special case

⁴The expectation is with respect to the random choice of P .

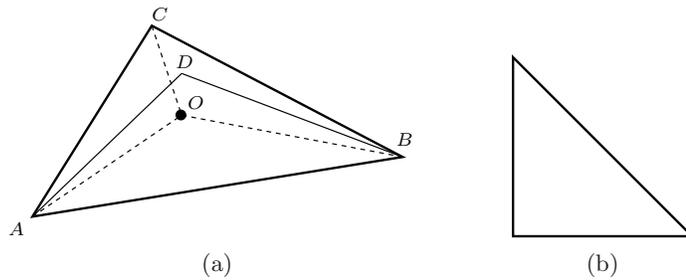


FIG. 5. (a) The canonization step. The triangle ABC is covered by three triangles, each of which contains the center O of the inscribed circle of ABC and has two edge orientations that are taken from a fixed set of $O(1/\alpha)$ directions. Only one of these triangles (ABD) is depicted in the figure. (b) A semicanonical right triangle after an appropriate affine transformation.

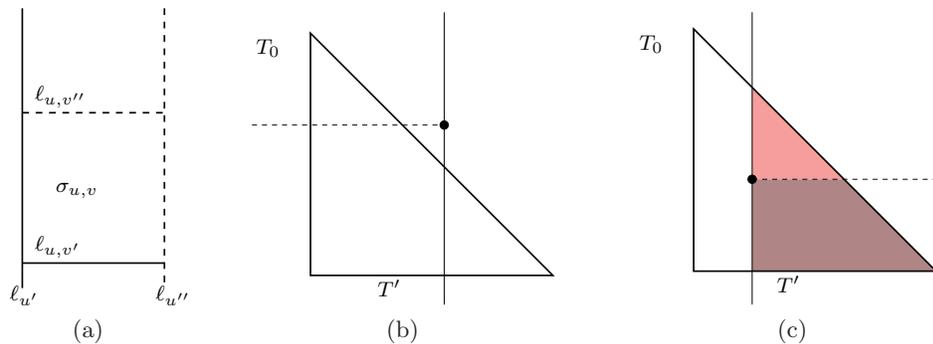


FIG. 6. (a) The “quadrant” $\sigma_{u,v}$ is defined by the line splitters $\ell_{u'}$, $\ell_{u,v'}$, but it is also bounded on the other sides by ancestor splitters $\ell_{u''}$ and $\ell_{u,v''}$. (b) The apex of $\sigma_{u,v}$, depicted by the black dot, misses the triangle T_0 when its heavy part T' is the left portion. (c) An anchored subtriangle T appears as a triangle (lightly shaded) homothetic to T_0 at the upper right quadrant or as a right-angle trapezoid (darkly shaded) at the lower right quadrant.

in the analysis, but for the time being, two levels suffice.) Following the notation of section 4, each node u of the primary tree is associated with the subset P_u of points that it represents, and a secondary (y -sorted) tree \mathcal{T}_u on P_u , and a node v of a secondary tree \mathcal{T}_u is associated with a corresponding subset $P_{u,v}$ of P_u . Each nonleaf node u of the primary tree stores a vertical line “splitter” ℓ_u , and a nonleaf node v of a secondary tree \mathcal{T}_u stores a horizontal line splitter $\ell_{u,v}$. For each such secondary node v of a tree \mathcal{T}_u , the lines $\ell_{u'}$ and $\ell_{u,v'}$, where u' is the parent of u in \mathcal{T} and v' is the parent of v in \mathcal{T}_u (as before, we handle only nodes for which u' and v' exist), define a quadrant $\sigma_{u,v}$, which is the intersection of two halfplanes bounded by $\ell_{u'}$ and $\ell_{u,v'}$ and containing $P_{u,v}$. (Technically, similar to the situation in section 4, $\sigma_{u,v}$ is a (possibly unbounded) rectangle, where the other vertical and horizontal edges of $\sigma_{u,v}$, if they exist, are portions of respective splitters $\ell_{u''}$, $\ell_{u,v''}$, where u'' is an appropriate ancestor of u' in \mathcal{T} and v'' is an appropriate ancestor of v' in \mathcal{T}_u ; see Figure 6(a).)

Let T_0 be a right triangle in our semicanonical family \mathcal{F} containing at least $\varepsilon n/3 = 44n/r$ points of P . We locate first the highest node u' in \mathcal{T} so that the line $\ell_{u'}$ meets T_0 , thus splitting it into two parts, where the right part is a triangle, homothetic to T_0 , and the left part is a right-angle trapezoid. Let T' be the part that contains at least $22n/r$ points of P , and let u be the corresponding child of u' so that its slab σ_u contains T' . We next locate the highest node v' in \mathcal{T}_u such that $\ell_{u,v'}$ meets

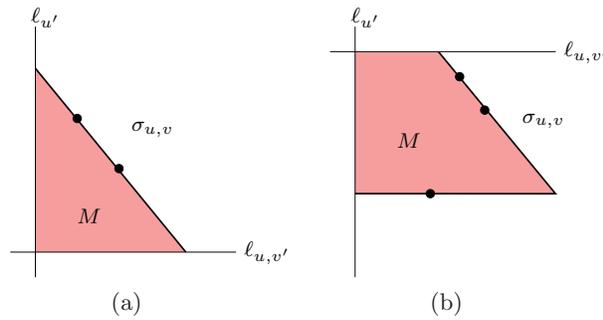


FIG. 7. A maximal anchored $R_{u,v}$ -empty (a) right triangle and (b) right trapezoid.

T' , splitting it into two subparts. We focus on the part T of T' that contains at least $11n/r$ points and denote by v the child of v' whose corresponding quadrant $\sigma_{u,v}$ contains T .

We observe that when T' is the right portion of T_0 , the apex o of $\sigma_{u,v}$ must be contained in T_0 , which is helpful for our analysis since then it implies that T is anchored at o in a sense similar to that in the preceding sections. However, this property does not necessarily hold when T' is the left portion of T_0 , in which case the analysis is more involved and some part of it requires an additional structure built on top of the tree \mathcal{T} —see Figure 6(b) and below. We thus begin with the analysis for the former case and then proceed to the more intricate latter case.

T' is the right portion of T_0 . The clipped region T is either (a) a triangle, homothetic to T_0 , whose right-angle vertex is the apex o of $\sigma_{u,v}$, or (b) a right-angle trapezoid, having o as its top left vertex, so that its bases are horizontal, its left side is vertical, and its right side is a portion of the hypotenuse of T_0 ; see Figure 6(c). In both cases we refer to T as being anchored at o . Note that in case (a) v is a right child of its parent, representing an upper quadrant, and that in case (b) v is a left child, representing a lower quadrant. Also, in both cases the slope of the slanted edge of T is negative, so in case (b) the slanted edge moves “away” from o , making the lower base of T longer than its upper base.

Recall that we have drawn a “global” random sample R of P . For each node v of each secondary tree \mathcal{T}_u , we put $R_{u,v} := R \cap \sigma_{u,v}$ and $r_{u,v} = |R_{u,v}|$. We next construct a family $\mathcal{M}_{u,v}$ of maximal anchored $R_{u,v}$ -empty regions, with the property that each anchored $R_{u,v}$ -empty region T (triangle or trapezoid within the right portion T' , as above), is covered by at most two regions in $\mathcal{M}_{u,v}$. (Below we will add more regions to $\mathcal{M}_{u,v}$, to handle other heavy subparts of T_0 .) Each region in $\mathcal{M}_{u,v}$ is either (a) an anchored $R_{u,v}$ -empty right triangle whose hypotenuse touches two points of $R_{u,v}$ (that is, it supports an edge of the convex hull of $R_{u,v}$), or (b) an anchored $R_{u,v}$ -empty right-angle trapezoid whose slanted side (has negative slope and) touches two points of $R_{u,v}$, and whose unanchored (lower) horizontal base passes through a point of $R_{u,v}$ (which might coincide with one of the two points lying on the slanted edge, i.e., be a vertex of the region), or else lies on the bottom side of the “quadrant” $\sigma_{u,v}$. In each of these cases, the region is clipped within $\sigma_{u,v}$ (when the actual rectangle $\sigma_{u,v}$ is delimited by more than two splitters). See Figure 7.

In case (a), we also include in $\mathcal{M}_{u,v}$ two axis-parallel rectangles M_1, M_2 anchored at o , so that (i) the right edge of M_1 passes through the leftmost point of $R_{u,v}$ and its top edge lies on the top side of $\sigma_{u,v}$ if it exists (otherwise M_1 extends to ∞),

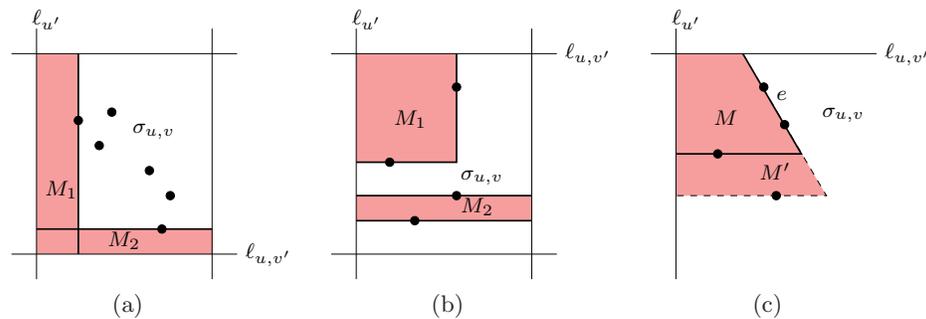


FIG. 8. Anchored maximal $R_{u,v}$ -empty rectangles (a) for upper right quadrants and (b) for lower right quadrants. (c) The larger right trapezoid M' cannot be empty if it shares its slanted edge e with M .

and (ii) the top edge of M_2 passes through the bottommost point of $R_{u,v}$ and its right edge lies on the right side of $\sigma_{u,v}$ if it exists (otherwise M_2 extends to ∞). See Figure 8(a). In case (b), we also include in $\mathcal{M}_{u,v}$ axis-parallel rectangles of the following two types: (i) rectangles that are anchored at o , with both right and bottom sides passing through a point of $R_{u,v}$; (ii) rectangles whose left and right sides lie, respectively, on the left and right sides of $\sigma_{u,v}$ (if the right side exists), and whose top and bottom sides pass through two respective points of $R_{u,v}$, necessarily consecutive in the y -order (including two extreme rectangles, above the highest point and below the lowest point). See Figure 8(b). Finally, if $R_{u,v}$ is empty, then $\mathcal{M}_{u,v}$ consists of the single region $\sigma_{u,v}$.

We next claim that $|\mathcal{M}_{u,v}| = O(r_{u,v} + 1)$. This is trivial when $R_{u,v} = \emptyset$, so assume that $R_{u,v}$ is nonempty. The claim is then obvious for regions of type (a), because their number is at most two plus the number of edges of the lower left convex hull of $R_{u,v}$. To bound the number of regions of type (b), sort the points of $R_{u,v}$ in decreasing y -order, and let the sorted sequence be $(q_1, q_2, \dots, q_{r_{u,v}})$. Put $R^{(j)} = \{q_1, \dots, q_{j-1}\}$ for $j = 1, \dots, r_{u,v}$. Let M be a region of type (b) whose lower horizontal base passes through q_j so that q_j is not a vertex of M . Then its slanted edge must contain an edge e of the (lower left) convex hull of $R^{(j)}$. Moreover, if such an M exists, then there cannot exist another region M' whose slanted edge contains e and whose lower base passes through any point q_k with $k > j$; see Figure 8(c). If q_j is the lower right vertex of M , the other point lying on the slanted edge belongs to $R^{(j)}$ and is uniquely determined. Hence the number of regions of type (b) (ignoring the extreme rectangular regions) is upper bounded by $r_{u,v}$ plus the overall number of distinct edges of the “incremental” convex hulls of $R^{(1)}, \dots, R^{(r_{u,v})}$. The latter number is $O(r_{u,v})$ because every newly added point q_j generates one new edge of the modified hull, possibly deleting several other edges from the hull. (Note that this is exactly the analysis of the classical “Graham scan” convex hull algorithm.) There are only two extreme rectangular ranges of type (a) in $\mathcal{M}_{u,v}$. The number of extreme rectangular ranges of type (b) is easily seen to be $O(r_{u,v} + 1)$, using a variant of the analysis in section 2.

The collection \mathcal{M} of canonical covering regions (as in section 3) includes so far the union of all the sets $\mathcal{M}_{u,v}$, over all primary nodes u and all nodes v of the respective secondary trees \mathcal{T}_u , for a total of $|\mathcal{M}| = O(|R| \log^2 r + r \log r)$ regions.

We also have the following promised property: Let T be the remaining portion of an initial triangle T_0 , and let u and v be the respective primary and secondary nodes

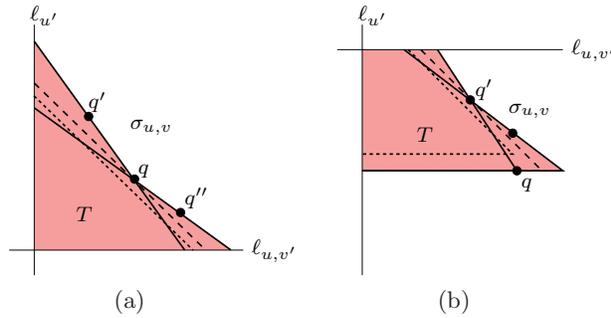


FIG. 9. The dotted edges are those of the original triangle or trapezoid T . The dashed edges are the slanted edges of appropriate expansions of the original T . Each such expansion is contained in the union of a pair of regions of $\mathcal{M}_{u,v}$.

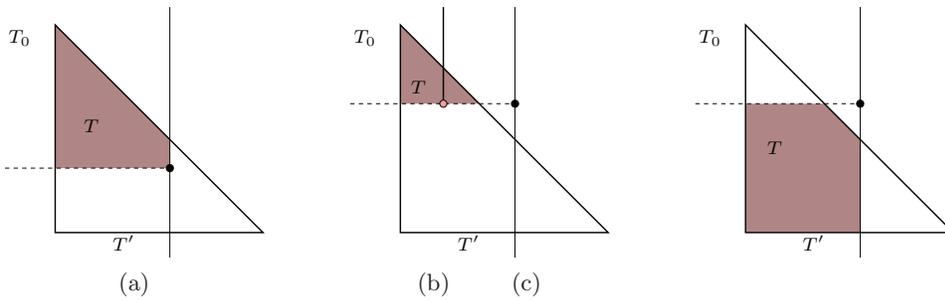


FIG. 10. (a) T is an anchored trapezoid. (b) The apex of the tertiary quadrant, depicted by the lightly shaded dot, hits the triangle T . (c) T is a truncated axis-parallel rectangle.

for which T is an anchored triangle or trapezoid within $\sigma_{u,v}$, as constructed above. Then, if T is $R_{u,v}$ -empty, it is contained in the union of at most two regions of $\mathcal{M}_{u,v}$.

Indeed, we may assume that $R_{u,v} \neq \emptyset$. Suppose first that T is a triangle. Translate the hypotenuse of T away from the apex o of $\sigma_{u,v}$, until it passes through a point q of $R_{u,v}$ (necessarily a hull vertex). Then rotate the new hypotenuse about q clockwise (resp., counterclockwise) until it meets a second point q' (resp., q'') of $R_{u,v}$ or becomes vertical (resp., horizontal). The two resulting triangles (or rectangles in the extreme cases) belong to $\mathcal{M}_{u,v}$, and their union covers T . See Figure 9(a).

Suppose next that T is a trapezoid. Expand T downward by sliding its bottom edge parallel to itself, while keeping the remaining bounding lines fixed, until its bottom edge hits some point $q = q_j$ of $R_{u,v}$ (that is, in the above notation, q is the j th highest point of $R_{u,v}$) or else reaches the lower boundary of $\sigma_{u,v}$. Then translate the slanted edge of the new trapezoid to the right until it hits a point q' of $R_{u,v}$ (more precisely, of $R^{(j)}$). Finally, rotate the new slanted edge about q' clockwise and counterclockwise until it meets a second point of $R^{(j)}$ or becomes vertical or horizontal; the clockwise rotation may end when it hits $q = q_j$. This yields two trapezoids (or rectangles) of $\mathcal{M}_{u,v}$ whose union covers T . See Figure 9(b). (Note that in both cases, the expansion of T may fall outside $\sigma_{u,v}$. This, however, does not violate our analysis, since in this case T is still contained in the union of at most two regions of $\mathcal{M}_{u,v}$, possibly clipped within $\sigma_{u,v}$.)

T' is the left portion of T_0 . We next consider the more involved case where T' is the left portion of T_0 .

First consider the case where T is the top portion of T' . Then T is a right-angle trapezoid, its bases are vertical, its bottom side is horizontal, and its upper side is a portion of the hypotenuse of T_0 . Moreover, in this case T has o as its bottom right vertex and is thus anchored at o ; see Figure 10(a). Hence, the current setting is a reflected rotation (by 90 degrees) of the scenario presented in the previous case (for lower right quadrants), and T can be covered by at most two maximal R -empty regions of a similar kind as those presented above, whose overall number is only linear in the size of $R_{u,v}$ within the quadrant $\sigma_{u,v}$. We add the new covering regions to our collection $\mathcal{M}_{u,v}$.

If T is a triangle, then it is no longer anchored at o . In this case, we extend the construction of the range-tree decomposition \mathcal{T} to include a third level, so that each tertiary tree $\mathcal{T}_{u,v}$ associated with a node v of a secondary tree \mathcal{T}_u stores the points of $P_{u,v}$ by their x -order. We construct each of these levels of \mathcal{T} down to nodes for which the size of their associated subset is n/r . Let w be the highest node in the tertiary tree $\mathcal{T}_{u,v}$ under consideration whose associated vertical line crosses T . This last step guarantees that each of the two resulting portions of T is now anchored at the apex o' of the respective new subquadrant $\sigma_{u,v,w}$, and the heavier of the two portions contains at least $5.5n/r$ points of P . See Figure 10(b). We thus face a scenario similar to the preceding case (where now the scene is reflected and rotated by 90 degrees), where we need to proceed with either an anchored triangle or an anchored trapezoid. Extending the definition of the regions placed so far in \mathcal{M} in an appropriate manner, we can conclude that T is covered by at most two maximal R -empty regions of the new kind, or, rather, $R_{u,v,w}$ -empty regions, where $R_{u,v,w} = R \cap \sigma_{u,v,w}$. We add these new covering regions (for the two subcases where T is a triangle or a trapezoid) to \mathcal{M} and note that the size of \mathcal{M} increases by a $\log r$ -factor to $O(|R| \log^3 r + r \log^2 r)$.

The more subtle case is when T is the bottom portion of T' . If o lies in T_0 , then T is a heavy axis-parallel rectangle (anchored at o) and is therefore stabbed by the initial net N_1 , so we may disregard this case.

Consider then the case where o misses T_0 . Then T is a truncated axis-parallel rectangle obtained from an axis-parallel rectangle (anchored at o) by removing a right-angle triangle, whose right angle is at o and which shares its hypotenuse with that of T_0 ; see Figure 10(c). We first trim T into a right-angle triangle with two vertices lying on the boundary of $\sigma_{u,v}$. For this we remove from T at most two axis-parallel rectangles, as illustrated in Figure 11(a), and recall that by our assumption neither of them contains more than n/r points of P , so the trimmed portion, which we continue to denote as T , is still heavy, containing at least $9n/r$ points.

Let β^+ , β^- denote the two rays emanating from o into $\sigma_{u,v}$ and forming angles of, say, 40° , with the top and right sides of $\sigma_{u,v}$, respectively. These rays partition $\sigma_{u,v}$ into three “sectors,” denoted, in their counterclockwise order around o , as $\sigma_{u,v}^+$, $\sigma_{u,v}^0$, and $\sigma_{u,v}^-$. The angle (40°) is chosen so as to guarantee the following two properties. (a) The right-angle vertex of T lies in $\sigma_{u,v}^0$. (b) Let a and b denote the respective intersection points of β^+ with the hypotenuse and the left side of T , and let c and d denote the respective intersection points of β^- with the hypotenuse and the bottom side of T . Then c lies above b and a lies to the right of d , as follows by simple trigonometric analysis. See Figure 11(b). This “monotonicity” property is crucial for the analysis—see below.⁵

The partition of $\sigma_{u,v}$ induces a partition of T into three corresponding portions,

⁵This is the only part of the analysis where we use the fact that the slope of the hypotenuse is close to -1 .

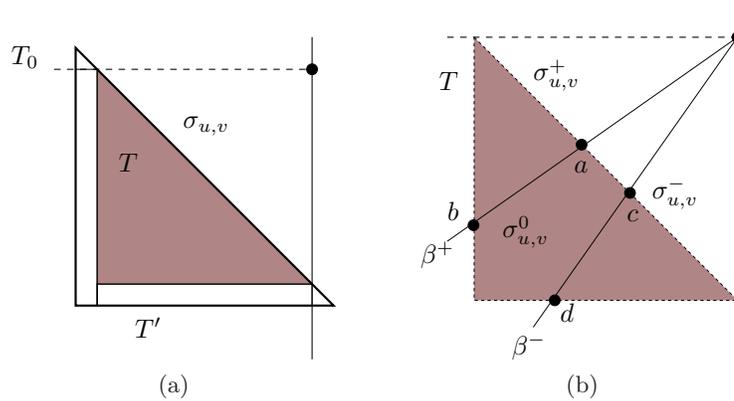


FIG. 11. (a) The trimmed triangle T is depicted by the shaded region. (b) The partition of $\sigma_{u,v}$ into three sectors.

referred to as its top, middle, and right portions. Since T is heavy, at least one of these portions contains at least $3n/r$ points. We consider only the cases where the heavy portion is the top or the middle portion. The right portion is handled by a symmetric argument.

A top heavy portion. So assume first that the heavy portion is $T^+ = T \cap \sigma_{u,v}^+$, and refer to Figure 12(a). Translate the “hypotenuse” of T^+ (the edge contained in the hypotenuse of T_0) toward o , while keeping its endpoints lying on the top boundary γ of $\sigma_{u,v}$ and on β^+ , respectively, until it hits a point $p \in R$. Then rotate its top portion clockwise around p and its bottom portion counterclockwise around p , until each of them hits a second point of R , which we denote, respectively, as p^+ and p^- , or becomes vertical or horizontal, respectively.⁶

Assume that both p^+ and p^- exist; the other cases are treated similarly and are, in fact, simpler. Let T^* be the region bounded by the left edge of T , by β^+ , by a portion of the top side of $\sigma_{u,v}$, and by the two segments $e^+ = pz^+$, $e^- = pz^-$, extending from p through p^+ and p^- to the top boundary of $\sigma_{u,v}$ and to β^+ , respectively; see Figure 12(a). Clearly, T^* contains T^+ .

We cover T^* by two regions, both R -empty. The top region Δ^+ is a right trapezoid whose bases lie on the top edge of $\sigma_{u,v}$ and on the horizontal line through p , whose slanted edge is e^+ and whose vertical edge passes through a point of R (or extends all the way to the left edge of $\sigma_{u,v}$ if it exists). The bottom region Δ^- is a quadrangle bounded by the horizontal line through p , by e^- , by β^+ , and by a fourth left vertical edge passing through a point of R (or else it extends all the way to the left edge of $\sigma_{u,v}$); see Figure 12(b).

It is easily checked that Δ^+ and Δ^- are uniquely determined by e^+ and e^- , respectively, for a fixed choice of R and for a fixed quadrant $\sigma_{u,v}$. We add all these regions to our canonical collection $\mathcal{M}_{u,v}$ (and eventually to \mathcal{M}).

We next argue that the number of these R -empty (or rather $R_{u,v}$ -empty) regions, within $\sigma_{u,v}$, is linear in $|R_{u,v}|$. That is, we claim that the number of edges e^+ , e^- that can arise in such a construction is linear in $|R_{u,v}|$. Indeed, consider, say, the edge e^+ , defined by the pair $p, p^+ \in R_{u,v}$. We claim that a fixed point p^+ can participate

⁶Note that rotating only the top and bottom portions separately guarantees that the point p^- at which the counterclockwise rotation stops cannot lie to the left of T .

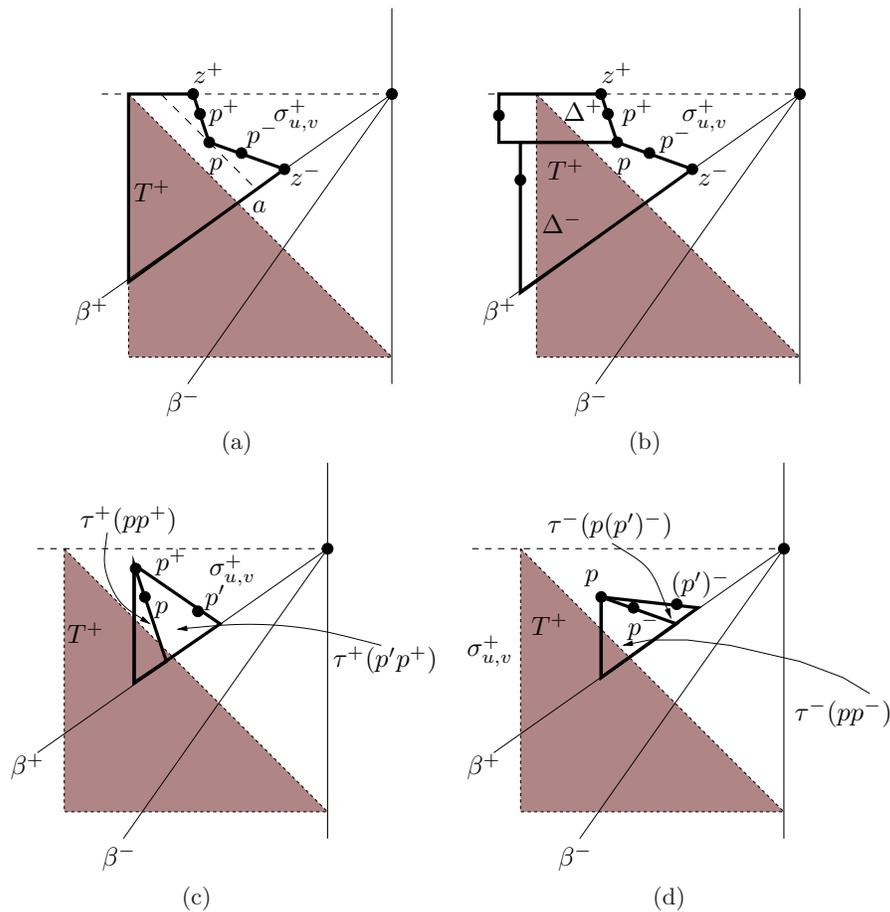


FIG. 12. (a) T^* , depicted by the bold polygon, contains T^+ . (b) The union of the two R -empty regions Δ^+ , Δ^- covers T^* . (c) Illustrating the argument that a given p^+ can be coupled with at most one point p . (d) Illustrating the argument that a given p can be coupled with at most one point p^- .

in at most one such pair. Indeed, consider the triangle $\tau^+(pp^+)$, bounded by the line containing e^+ , by β^+ , and by the vertical line through p^+ ; see Figure 12(c). The construction guarantees that this triangle is $R_{u,v}$ -empty. Hence, if p^+ participated in another edge $(e')^+$, with another lower point $p' \in R_{u,v}$ which lies counterclockwise to p (about p^+), then p would have to lie in the corresponding triangle $\tau^+(p'p^+)$, contradicting the fact that this triangle is $R_{u,v}$ -empty; see Figure 12(c). Similarly, a fixed point p can participate in at most one pair p, p^- . Indeed, consider the triangle $\tau^-(pp^-)$, bounded by the line containing e^- , by β^+ , and by the vertical line through p ; see Figure 12(d). The construction guarantees that this triangle is $R_{u,v}$ -empty. Hence, if p participated in another edge $(e')^-$, with another lower point $(p')^- \in R_{u,v}$ which lies counterclockwise to p^- , then p^- would have to lie in the corresponding triangle $\tau^-(p(p')^-)$, contradicting the fact that this triangle is $R_{u,v}$ -empty.

A middle heavy portion. Assume next that the heavy portion is $T^0 = T \cap \sigma_{u,v}^0$, which is a pentagon bounded by portions of β^+ , β^- , the left and bottom edges of T , and its hypotenuse.

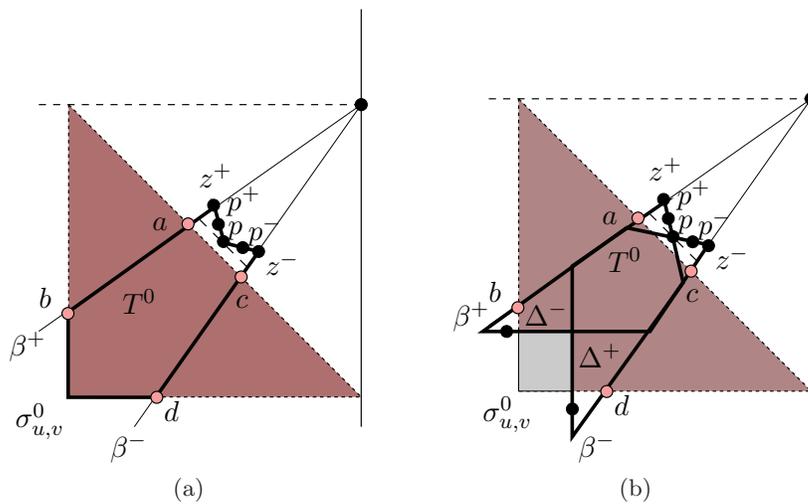


FIG. 13. (a) T^* , depicted by the bold polygon, contains T^0 . (b) The union of the regions Δ^+ , Δ^- , depicted by the two respective bold polygons, covers T^* , except possibly for an axis-parallel rectangle, drawn lightly shaded.

Refer to Figure 13, and translate the “hypotenuse” of T^0 (the edge lying on the hypotenuse of T_0) toward o , while keeping its endpoints on β^+ and on β^- , respectively, until it hits a point $p \in R$. Then rotate its top portion clockwise around p and its bottom portion counterclockwise around p , until each of them hits a second point of R , which we denote, respectively, as p^+ and p^- , or becomes vertical or horizontal, respectively.

Assume that both p^+ and p^- exist; as above, the other cases are treated similarly and are in fact simpler. Let T^* be the region bounded by (portions of) β^+ , β^- , the left and bottom edges of T , and the two segments $e^+ = pz^+$, $e^- = pz^-$, extending, respectively, from p through p^+ and p^- to β^+ and to β^- ; see Figure 13(a). Clearly, T^* contains T^0 .

We next argue that the number of edges e^+ , e^- that can arise in the above construction is linear in $|R_{u,v}|$. The argument more or less follows that given above for T^+ . That is, consider, say, the edge e^+ , defined by the pair $p, p^+ \in R$, and observe, as above, that a fixed point p^+ can participate in at most one such pair. Similarly, a fixed point p can participate in at most one pair p, p^- .

We cover T^* by the following two R -empty regions, plus an axis-parallel rectangle. We first observe that, due to the monotonicity property stated above, c lies above b , and a lies to the right of d , and it is thus guaranteed that p, p^+, p^- lie above b and to the right of d . Let E^+ (resp., E^-) denote the segment between β^+ and β^- passing through p and p^+ (resp., through p and p^-). By construction, E^+ lies fully to the right of a and thus also to the right of d . Let Δ^+ be the $R_{u,v}$ -empty quadrangle bounded by β^+ , β^- , E^+ , and a left vertical edge which passes through a point of $R_{u,v}$ (or else extending all the way to the left side of $\sigma_{u,v}$). Similarly, E^- lies fully above c and thus also above b . Let Δ^- be the $R_{u,v}$ -empty quadrangle bounded by β^+ , β^- , E^- , and a bottom horizontal edge which passes through a point of $R_{u,v}$ (or else extending all the way to the bottom side of $\sigma_{u,v}$). Note that Δ^+ (resp., Δ^-) is uniquely determined by E^+ (resp., by E^-).

Either $\Delta^+ \cup \Delta^-$ fully cover T^0 , or else they leave out an axis-parallel rectangle,

whose top edge lies below b and whose right edge lies to the left of d . By assumption, this rectangle contains fewer than n/r points of P , and since T^0 contains at least $3n/r$ points, it follows that each of Δ^+ , Δ^- must contain at least n/r points of P . See Figure 13(b). We add these regions to our collection $\mathcal{M}_{u,v}$ (and, consequently, to \mathcal{M}).

Applying Theorem 3.2. The preceding analysis fits the abstract framework of Theorem 3.2. The set \mathcal{R}_1 is the collection of all the covering regions that we have used in the various subcases, including triangle, trapezoids, rectangles, and the regions constructed in the second part of the analysis (in which T' is the left portion of T_0), without the requirement that they be R -empty and maximal. We note that each of these regions is defined by at most three points of P . For a given random sample R , the set \mathcal{R}'_1 is the corresponding set \mathcal{M} that we have used above. The set \mathcal{R}_0 is the collection consisting of all the portions T (of original fat triangles T_0) of the various kinds considered above, clipped to within the respective regions, which are either quadrants $\sigma_{u,v}$, over all pairs u, v , portions of lower left quadrants delimited by the corresponding rays β^+ , β^- , or portions of upper left secondary quadrants split by a tertiary vertical line. The preceding analysis shows that each $\varepsilon n/3$ -heavy and R -empty triangle of \mathcal{F} (recall that in the canonization step each original triangle has been covered by three semicanonical triangles) contains a $(3n/r)$ -heavy range of \mathcal{R}_0 , and each range of \mathcal{R}_0 is covered by at most two regions of \mathcal{M} (and possibly by a third axis-parallel rectangle). We recall that the size of \mathcal{M} is $O(|R| \log^3 r + r \log^2 r)$, so its expected value is $O(s \log^3 r)$. Hence, we are within the abstract framework of section 3, so we can conclude the following.

THEOREM 6.1. *For any set P of n points in the plane, any fixed constant parameter $\alpha > 0$, and a parameter $\varepsilon > 0$, there exists an ε -net of P , for α -fat triangles, of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$, where the constant of proportionality depends on α .*

Constructing the ε -net. We construct an ε -net of this size using an easy variant of the algorithms presented in sections 2 and 4. For each of the $O(1/\alpha^3)$ semicanonical families, we apply an affine transformation to the plane and to P , which turns the two fixed edge directions into the coordinate directions; by a slight abuse of notation we continue to refer to the transformed set as P . Consider one such family. We construct the two-level range tree \mathcal{T} over the points of P , using $O(n \log n)$ time and storage. For the first case of the analysis, where T' is the right portion of T_0 , we enumerate the maximal anchored $R_{u,v}$ -empty regions M in each canonical quadrant $\sigma_{u,v}$ by tracking the edges appearing on the “incremental” convex hull of the points in $R_{u,v}$ for lower right quadrants, or by just enumerating the edges of the lower left hull of $R_{u,v}$ for upper right quadrants. We can produce these regions in time $O(r_{u,v} \log r_{u,v})$, although we still need to test which of them are $R_{u,v}$ -empty. For simplicity, we perform this step by brute force. This takes a total of $O(1 + r_{u,v}^2)$ time per node, so the overall cost of producing these canonical empty regions is $O(s^2)$, as is easily checked (in this case the time bounds constitute a geometric sequence over the various levels of the tree); we assume s to be sufficiently small (specifically, $s = o(n^{1/2})$) to make this bound at most linear in n . Finding the degenerate canonical empty rectangles can be done by applying enumeration algorithms similar to those in section 2. For the latter case of the analysis (where T' is the left portion of T_0), we use a construction similar to that used above for the case where T is the top portion of T' , with an appropriate reflection and rotation, applying it at the tertiary level of the tree. When T is the bottom portion of T' , we partition each quadrant at the secondary level into three sectors as above, and in each of them we generate the maximal empty regions Δ^+ ,

Δ^- in a brute-force manner. That is, for each point p_1 in, say, $\sigma_{u,v}^+$, we search in linear time for the first point p_2 of $R_{u,v}$ in that sector which lies counterclockwise about p_1 to the downward vertical ray emanating from p_1 . For each such pair of points (p_1, p_2) , we compute the regions Δ^+ , Δ^- that they induce (where the bottom edge of Δ_1 passes through p_2 and the top edge of Δ_2 passes through p_1) in a brute-force manner, by examining each of the remaining points of $R_{u,v}$ as the potential left delimiter of either region. We use similar constructions to produce the corresponding regions in $\sigma_{u,v}^0$, $\sigma_{u,v}^-$. As above, since the number of such regions in a fixed sector is only linear in the number of points, the overall running time, over all sectors and all nodes in the tree, is $O(s^2)$.

We next compute the weight factor t_M of each of the $O(s \log^3 r)$ maximal empty regions $M \in \mathcal{M}$. For this, we prepare an appropriate version of a triangle range counting structure in the plane, which uses linear storage and $O(n \log n)$ preprocessing time and answers queries in time $O(n^{1/2} \text{polylog } n)$ [Mat92a]. The overall cost of answering $O(s \log^3 r)$ queries, including preprocessing, is thus $O(n \log n + sn^{1/2} \log^3 r \text{polylog } n)$, which is only $O(n \log n)$, for $s = O(n^{1/2} / \log^c n)$, for a suitably large absolute constant $c > 0$. We then proceed in a similar manner as that described for the previous algorithms in sections 2 and 4. Omitting any further details, we obtain that the overall expected running time of the algorithm is $O(n \log n)$, with a constant of proportionality that depends on α (for $s = O(n^{1/2} / \log^c n)$, for some a suitably large positive absolute constant c).

7. Improved bounds for ε -nets for dual range spaces. In this section we observe that the technique developed in this paper can be adapted to the scenarios considered by Clarkson and Varadarajan [CV07] and yields improved bounds for the size of ε -nets in many of the cases considered there.

Rephrasing the notation used in the introduction, we consider the dual range space $\Xi = (\mathcal{C}, \mathcal{Q})$, where the ground set \mathcal{C} is a collection of geometric regions in \mathbf{R}^d , and each range in \mathcal{Q} is of the form $Q_x = \{C \in \mathcal{C} \mid x \in C\}$ for some $x \in \mathbf{R}^d$. Let us consider a finite subcollection \mathcal{C}' of any m regions of \mathcal{C} , and let W' denote the complement of their union. We assume that W' can be decomposed into *cells* of some simple shape. That is, each cell M in the decomposition is a (possibly unbounded) portion of W' that is *defined* by a set $D(M)$ of $O(1)$ regions of \mathcal{C}' , in the sense that it appears in the decomposition of the complement of the union of just those $O(1)$ regions (in particular, the cells of the decomposition do not necessarily have the same shape as the regions of \mathcal{C}). In the terminology used in the probabilistic framework of Clarkson and Shor [CS89], we call $D(M)$ the *defining set* of M . We also associate with M a *killing set* $K(M)$, consisting of all regions in \mathcal{C}' that intersect M . In many geometric range spaces of this kind, the cells are those generated by the *vertical decomposition* of the complement of the union [SA95], although there exist other types of decompositions for various special classes of regions; see, e.g., [AMS98, Cla87, CS89] for a description of this (standard) set-up.

Clarkson and Varadarajan [CV07] further assume that the number of cells M in the complement of the union is at most $m\varphi(m)$, where $\varphi(m)$ is some slowly increasing sublinear function (recall the definitions in section 3). Under these assumptions, Clarkson and Varadarajan show that the range space Ξ admits ε -nets of size $O(\frac{1}{\varepsilon}\varphi(\frac{1}{\varepsilon}))$. Thus, if $\varphi(m) = o(\log m)$, the resulting nets are smaller than the standard bound $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ of [HW87].

In this section we obtain the following improvement.⁷

THEOREM 7.1. *Under the assumptions made above, the range space Ξ admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \varphi\left(\frac{1}{\varepsilon}\right)\right)$ for any $0 < \varepsilon \leq 1$.*

Remarks. (1) The bound in the theorem improves upon the general bound $O\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ when $\varphi(m) = 2^{o(\log m)}$, thus extending the applicability of this technique beyond the “effective range” $\varphi(m) = o(\log m)$, where the original technique of [CV07] yields an improvement.

(2) The probabilistic model that we choose (similarly to the model in section 2) yields a considerably simpler analysis than that in [CV07].

Proof. We follow some of the ideas of the general approach of section 3 and adapt them to the current scenario. Here we have a finite subcollection of n elements of \mathcal{C} , which, for simplicity, we continue to denote by \mathcal{C} . We put $r := 1/\varepsilon$, $s := cr \log \varphi(r)$, and $\pi := s/n$, where $c > 1$ is a constant. We draw a random sample R of regions of \mathcal{C} , picking each region, independently, with probability π . We form the union \mathcal{U} of R and decompose its complement into at most $|R|\varphi(|R|)$ simply shaped regions, each determined by $O(1)$ sets of R ; as above, we refer to the regions which form the decomposition as “cells.” We define the weight factor t_M of a cell M to be $s|\mathcal{C}_M|/n$, where \mathcal{C}_M is the subcollection of those regions of \mathcal{C} which meet M . By the standard ε -net theory [HW87], or, alternatively, by the Clarkson–Shor technique [Cla87, CS89], it follows that, with high probability, we have⁸ $|\mathcal{C}_M| = O\left(\frac{n}{s} \log s\right)$ for each cell M , and, in an informal and imprecise sense, the expected size of \mathcal{C}_M , for a cell M , is only $O(n/s)$.

As above, we take each “heavy” cell M , with $t_M \geq c \log \varphi(r)$, and use the standard theory of ε -nets to deduce that there exists a $(1/t_M)$ -net N_M for \mathcal{C}_M , whose size is $O(t_M \log t_M)$. We output the union of R with all the sets N_M , over all heavy cells M , as the desired $(1/r)$ -net (that is, ε -net) N .

Adapting the argument in section 3, it is straightforward to verify that N is indeed an ε -net. Recall that in this dual context an ε -net is a subset of regions covering all points that are contained in at least an ε -fraction of the regions. To bound the expected size of N , we follow the same analysis as in section 3. That is, we apply the Exponential Decay Lemma in this context, where the defining and killing sets of a cell M are defined as above. In essentially all cases considered in [CV07] and below, the axioms assumed in [AMS98], or their simplified version used in section 3, hold. We denote by $\text{CT}(R)$ the set of all cells appearing in the decomposition of the complement of the union of a subset R of \mathcal{C} , and by $\text{CT}_t(R)$ the subset of $\text{CT}(R)$ consisting of those cells with weight factor at least t .

It thus follows that the Exponential Decay Lemma is applicable in this scenario as well. The rest of the analysis is essentially identical to that in Theorem 3.2. We thus conclude that the overall expected size of the subsamples “within” each heavy cell of the complement of the union is sublinear in r , so the expected size of N is dominated by that of R ; thus it is $O(r \log \varphi(r))$. \square

Several special cases. Theorem 7.1 immediately implies improved bounds on the size of ε -nets for dual range spaces of several classes of regions and points, for which the union complexity (or, rather, the complexity of the decomposition of its complement)

⁷Of course, it is an improvement only when φ is $\omega(1)$; otherwise, the bound is $O(1/\varepsilon)$, as already follows from [CV07].

⁸Normally, for these bounds to hold, one needs to consider only those regions of \mathcal{C} which *cross* (i.e., intersect but do not fully contain) M . However, in our case we do not need this distinction: Since each cell M is disjoint from all regions of R , the above analysis also applies to regions of \mathcal{C} that fully contain M .

is known to be nearly linear. We first present some of the standard families with this property and state their union complexity. Since these are *families* of planar regions, the following bounds also apply, with some care, for the complexity of the decomposition of the complement of their union. (We consider only families for which the known bound is superlinear; there is no improvement when the union complexity is linear.)

α -fat triangles (Figure 14(a)). Recall that a triangle is α -fat if each of its angles is at least α . The complexity of the union of n such triangles is $O(n \log \log n)$, where the constant of proportionality depends on the fatness factor α [MPSSW94, PT02].

Locally γ -fat objects (Figure 14(b)). These objects were recently introduced by de Berg [dB08, dB10]. Given a parameter $0 < \gamma \leq 1$, an object o is *locally γ -fat* if, for any disk D whose center lies in o , such that D does not fully contain o in its interior, we have $\text{area}(D \cap o) \geq \gamma \cdot \text{area}(D)$, where $D \cap o$ is the connected component of $D \cap o$ that contains the center of D . We also assume that the boundary of each of the given objects has only $O(1)$ locally x -extreme points, and that the boundaries of any pair of objects intersect in at most s points, for some constant s . Both conditions, for example, are satisfied if we assume that each γ -fat object has constant description complexity. It is then shown in [dB10] that the combinatorial complexity of the union of n such objects is $O(\lambda_{s+2}(n) \log n)$, with a constant of proportionality that depends on γ , where $\lambda_t(q)$ is the nearly linear maximum length of Davenport–Schinzel sequences of order t on q symbols (see [SA95]). When the objects have roughly the same size (i.e., the ratio of the diameters of any pair of objects is bounded by some constant), the complexity of the union decreases to $O(\lambda_{s+2}(n))$. Locally γ -fat objects are a generalization of several other previously studied classes of “fat” objects [Ef05, EK99, ES00].

Semiunbounded pseudotrapezoids (Figure 14(c)). Here each object is a region of one of the forms

$$\begin{aligned}\tau_{x_1, x_2, f}^- &= \{(x, y) \mid x_1 \leq x \leq x_2, y \leq f(x)\} \text{ or} \\ \tau_{x_1, x_2, f}^+ &= \{(x, y) \mid x_1 \leq x \leq x_2, y \geq f(x)\},\end{aligned}$$

where f is a continuous function. We assume that the graphs of any pair of these functions intersect in at most s points for some constant s . In this case the complexity of the union of any n such objects is $O(\lambda_{s+2}(n))$; see, e.g., [SA95]. If the objects are *pseudohalfplanes*, that is, $x_1 = -\infty$ and $x_2 = +\infty$ for each object, the bound on the union complexity slightly improves to $O(\lambda_s(n))$.

Jordan arcs with three intersections per pair (Figure 14(d)). Each object is bounded by some Jordan arc which starts and ends on the x -axis but otherwise lies above it, and by the portion of the x -axis between these endpoints, and each pair of the bounding Jordan arcs intersect at most three times. In this case the complexity of the union of any n such objects is $O(\lambda_3(n)) = O(n\alpha(n))$, where $\alpha(\cdot)$ is the (extremely slowly growing) inverse Ackermann function; see [EGH*89]. We also assume that the boundary of each object has only $O(1)$ locally x -extreme points.

Recall that the actual condition is about the complexity of a decomposition of the complement of the union, rather than just the complexity of the union itself. However, since we are dealing with planar objects of the above kind, the standard vertical decomposition technique (see, e.g., [SA95]) yields a decomposition whose complexity is proportional to that of the union, so the above bounds hold for the

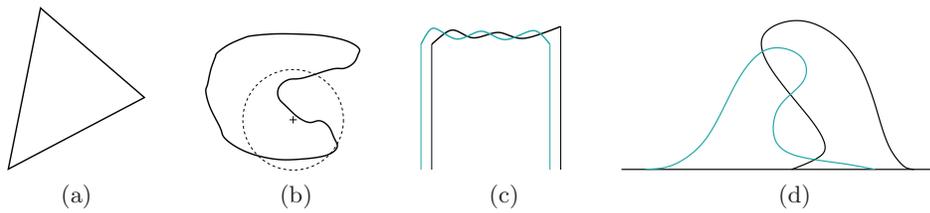


FIG. 14. The types of regions considered in this section: (a) An α -fat triangle; (b) a locally γ -fat region; (c) semiunbounded pseudotrapezoids; and (d) regions bounded by Jordan arcs with three intersections per pair.

decomposition as well.⁹

As noted by Clarkson and Varadarajan [CV07], their general approach implies that any dual range space of α -fat triangles and points admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$. Similarly, any dual range space of locally γ -fat objects and points, where the objects have roughly the same size, and each pair of object boundaries intersect in at most s points, admits an ε -net of size¹⁰ $O\left(\lambda_{s+2}\left(\frac{1}{\varepsilon}\right)\right)$. When the objects are bounded by Jordan arcs with three intersections per pair, as defined above, the size of the net becomes $O\left(\frac{1}{\varepsilon} \alpha\left(\frac{1}{\varepsilon}\right)\right)$.

Using Theorem 7.1 we can improve each of these bounds of [CV07], and also extend the bound for the case of locally γ -fat objects of *arbitrary* sizes (a case that cannot be treated by the original technique of [CV07]). That is, we have the following corollary.

COROLLARY 7.2. (a) *Any dual range space of α -fat triangles and points in the plane admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \log \log \frac{1}{\varepsilon}\right)$ for any $0 < \varepsilon \leq 1$.*

(b) *Consider a dual range space of locally γ -fat objects of arbitrary sizes in the plane and points, so that the boundary of each of the given objects has only $O(1)$ locally x -extreme points, and any pair of these boundaries meets in at most s points for s constant. Then any such dual range space admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon}\right)$ for any $0 < \varepsilon \leq 1$. When these objects have roughly the same size, the bound improves to $O\left(\frac{1}{\varepsilon} \log \beta_{s+2}\left(\frac{1}{\varepsilon}\right)\right)$, where $\beta_t(1/\varepsilon) = \varepsilon \lambda_t(1/\varepsilon)$.*

(c) *Consider a dual range space of semiunbounded pseudotrapezoids and points in the plane, where, for any pair of trapezoids, the graphs of their bounding functions intersect in at most s points for some constant s . Then any such dual range space admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \beta_{s+2}\left(\frac{1}{\varepsilon}\right)\right)$ for any $0 < \varepsilon \leq 1$. When the pseudotrapezoids are pseudohalfplanes, the bound improves to $O\left(\frac{1}{\varepsilon} \log \beta_s\left(\frac{1}{\varepsilon}\right)\right)$.*

(d) *Consider a dual range space of objects and points, where each object is bounded by a Jordan arc which starts and ends on the x -axis and by the portion of the x -axis between these endpoints. Each bounding Jordan arc has only $O(1)$ locally x -extreme points, and each pair of these arcs intersects at most three times. Then any such dual range space admits an ε -net of size $O\left(\frac{1}{\varepsilon} \log \alpha\left(\frac{1}{\varepsilon}\right)\right)$ for any $0 < \varepsilon \leq 1$.*

Remark. Applying the known upper bounds on the quantities $\beta_s(n)$ (see [ASS89,

⁹This is why we need to assume that no object boundary “wiggles” too much.

¹⁰In fact, Clarkson and Varadarajan [CV07] applied their technique to the more restricted class of (α, β) -covered objects of roughly equal size (see [Ef05] for the definition and the union complexity bound) and obtained a similar bound; the same technique applies to the more general class of locally γ -fat objects.

Niv09]), we have

$$\log \beta_s(n) = \begin{cases} O(\alpha^{\lfloor (s-2)/2 \rfloor}(n)), & s \geq 2 \text{ even,} \\ O(\alpha^{\lfloor (s-2)/2 \rfloor}(n) \log \alpha(n)), & s \geq 3 \text{ odd,} \end{cases}$$

while $\beta_1(\cdot)$ and $\beta_2(\cdot)$ are constants.

In closing, we note that, although the technique, as laid out at the beginning of this section, can be applied in principle to dual range spaces in any dimension, we have managed to apply it only to planar dual range spaces. The reason is the scarcity of classes of regions in higher dimensions with linear or, rather, near-linear bounds on the complexity of the decomposition of the complement of their union. It would of course be interesting to find such classes, and to apply our new technique to them.

8. Improved approximation factors for geometric SET COVER and HITTING SET problems. In this section we plug the improved bounds on the size of ε -nets, as derived in the preceding sections, into the machinery of Even, Rawitz, and Shahar [ERS05] to obtain improved approximation factors for the corresponding SET COVER or HITTING SET problems.

We first briefly recall this technique. For simplicity, we only review the HITTING SET variant. We are given a range space (P, \mathcal{R}) , and the goal is to find a small subset H of P which meets every range in \mathcal{R} . The technique of Even, Rawitz, and Shahar [ERS05] is based on *LP-relaxation*, and can be interpreted as a simplification for the preceding technique of Brönnimann and Goodrich [BG95]. The technique assumes the availability of two black-box routines: (i) a routine for approximately solving a linear program of an instance of the “fractional HITTING SET problem” (see below); (ii) an ε -net finder, which is a procedure that, given any positive weight function w on X and $\varepsilon > 0$, constructs a *weighted ε -net* N for (X, \mathcal{R}) in the sense that N hits each range whose weight is at least $\varepsilon w(X)$, with the weight of a subset of X being the sum of weights of its elements.

We now proceed as follows. Put $P = \{p_1, \dots, p_n\}$ and $\mathcal{R} = \{R_1, \dots, R_m\}$. We first solve the *fractional HITTING SET problem* defined by the following LP-system:

$$\begin{aligned} \min \quad & \sum_{i=1}^n x_i \\ \text{s.t.} \quad & \sum_{p_i \in R_j} x_i \geq 1, \quad j = 1, \dots, m, \\ & x_i \geq 0, \quad i = 1, \dots, n. \end{aligned}$$

Let OPT be the value achieved at the minimum. We then normalize the above LP-system. Putting $z_i = \frac{x_i}{\text{OPT}}$ for each $i = 1, \dots, n$ and $\varepsilon = \frac{1}{\text{OPT}}$, we obtain the system

$$\begin{aligned} \max \quad & \varepsilon \\ \text{s.t.} \quad & \sum_{i=1}^n z_i = 1, \\ & \sum_{p_i \in R_j} z_i \geq \varepsilon, \quad j = 1, \dots, m, \\ & z_i \geq 0, \quad i = 1, \dots, n. \end{aligned} \tag{1}$$

The algorithm then consists of the following two steps: (i) Solve (1). Let $Z^* := \{z_1^*, \dots, z_n^*\}$ be the resulting weights and ε^* the corresponding value of ε . (ii) Invoke

a net-finder for the set system (P, \mathcal{R}) with respect to the weights given by Z^* and the parameter ε^* , which outputs a weighted ε^* -net H . The key observation of Even, Rawitz, and Shahar [ERS05] is that H is the actual hitting set that we seek. Indeed, since the sum of the weights of the points in each range exceeds ε^* , it follows by definition that each range has to be stabbed by H .

Even, Rawitz, and Shahar [ERS05] proposed several polynomial-time algorithms to solve step (i) (some of which rely on approximations, as the algorithms of Young [You95] and Plotkin, Shmoys, and Tardos [PST95]; see [ERS05] for details).

We next apply this technique to our three main (primal) range spaces, consisting of points and axis-parallel rectangles in the plane, of points and axis-parallel boxes in \mathbf{R}^3 , and of points and α -fat triangles in the plane. We have presented in sections 2, 4, and 6 algorithms that construct an ε -net for these cases in nearly linear time, and it is straightforward to generalize these algorithms to the weighted case within the same asymptotic time bound (using, e.g., the rounding technique of Matoušek [Mat95]). We thus obtain the following corollary.

COROLLARY 8.1. *There exists a randomized, expected polynomial-time algorithm that, given a set \mathcal{Q} of m axis-parallel rectangles and a set P of n points in the plane that hit \mathcal{Q} , computes a subset $H \subseteq P$ of $O(\text{OPT} \log \log \text{OPT})$ points that hit \mathcal{Q} , where OPT is the size of the smallest such set. The algorithm can be extended to the case of axis-parallel boxes and points in 3-space, and α -fat triangles and planar point sets, yielding similar approximation factors in both cases.*

Using the above machinery, we also obtain polynomial-time approximation algorithms for the SET COVER problems associated with the dual range spaces considered in section 7. As shown in [CV07, Theorem 2.3], the ε -net can be constructed in time that is polynomial in the size of the ground set and in $1/\varepsilon$. (In cases (b)–(d) below we also assume that several basic operations on a constant number of the input regions can be performed in constant time; these include finding the intersection points of the boundaries of a pair of objects and testing whether a given point lies in an object.)

COROLLARY 8.2. (a) *There exists a randomized, expected polynomial-time algorithm that, given a set P of n points in the plane and a set T of α -fat triangles that cover P , computes a set cover $T' \subseteq T$ for P of size $O(\text{OPT} \log \log \log \text{OPT})$, where OPT is the size of the smallest such set.*

(b) *There exists a randomized, expected polynomial-time algorithm that, given a set P of n points in the plane and a set T of locally γ -fat objects of arbitrary sizes that cover P , so that the boundary of each of the given objects has only $O(1)$ locally x -extreme points, and each pair of these boundaries intersects in at most s points for some constant s and computes a set cover $T' \subseteq T$ for P of size $O(\text{OPT} \log \log \text{OPT})$, where OPT is the size of the smallest such set. When the elements of T are roughly the same size (i.e., with diameters differing by at most a constant factor), the size of the set cover improves to $O(\text{OPT} \log \beta_{s+2}(\text{OPT}))$.*

(c) *There exists a randomized, expected polynomial-time algorithm that, given a set P of n points in the plane and a set T of semiunbounded pseudotrapezoids that cover P , bounded by x -monotone curves, each pair of which meet at most s times, computes a set cover $T' \subseteq T$ for P of size $O(\text{OPT} \log \beta_{s+2}(\text{OPT}))$, where OPT is the size of the smallest such set; the bound slightly improves to $O(\text{OPT} \log \beta_s(\text{OPT}))$, when the input regions are pseudohalfplanes.*

(d) *There exists a randomized, expected polynomial-time algorithm that, given a set P of n points in the plane and a set T of objects that cover P , each of which is bounded by some Jordan arc which starts and ends on the x -axis and by the portion*

of the x -axis between these endpoints, so that each bounding Jordan arc has only $O(1)$ locally x -extreme points, and each pair of these arcs intersects at most three times, computes a set cover $T' \subseteq T$ for P of size $O(\text{OPT} \log \alpha(\text{OPT}))$, where OPT is the size of the smallest such set.

The results of this section are summarized in Tables 1–2.

9. Concluding remarks and open problems. In this paper we achieved significant progress on the problem of bounding the size of ε -nets for several set systems, both in the primal and the dual (geometric) settings. We believe that the oversampling approach presented in this paper is of independent interest, and that it will find additional applications on related problems. Very recently, Chekuri, Clarkson, and Har-Peled [CCH09] have used this approach to improve the approximation factors on the *set multicover* problem for points in the plane and each of the families of planar regions considered in section 7. In problems of this kind, we wish to find a minimum cardinality subset of the given regions such that each input point p is covered by at least $d(p)$ regions, where $d(p)$ is a integer demand (requirement) for p . In this case Chekuri, Clarkson, and Har-Peled achieve in polynomial time approximation factors similar to those stated in Corollary 8.2 for the original SET COVER problem (that is, with unit demands).

We conclude the paper by stating several open problems raised by our study (several other open problems are mentioned throughout the paper).

(i) One may consider the dual version of the main problem that we have studied. Namely, we are given a collection \mathcal{C} of n axis-parallel rectangles, and each range is the subset of \mathcal{C} stabbed by some point in the plane. Here too the goal is to show the existence of a small-size ε -net, which is a small-size subset $\mathcal{C}' \subseteq \mathcal{C}$ whose union contains all the “deep” points (i.e., points contained in at least εn rectangles of \mathcal{C}). So far we do not know how to apply our method to this dual set-up. We note that Brönnimann and Lenchner, in their conference paper [BL04], claim, without a proof, the existence of ε -nets for this dual range space, of size $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$.

(ii) Another challenging open problem is to extend our machinery for axis-parallel boxes to dimensions $d \geq 4$. The anchoring trick used for $d = 3$ fails, because the number of maximal R -empty orthants in d -space can be $\Theta(|R|^{\lfloor d/2 \rfloor})$ in the worst case [KRSV08], and the challenge is to prune away most of these orthants and remain only with a nearly linear number of them. The more modest goal of constructing a *weak ε -net* in this setting (a weak ε -net is defined analogously to a (strong) ε -net that was considered in this paper, except that its elements are not constrained to be taken from the input set) was recently met in a follow-up study [Ezr09], where the bound is shown to be $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$ in any dimension d (with a constant of proportionality that depends on d). Another goal is to construct weak ε -nets of size $o(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$ for the (primal) range spaces that we have studied in this paper, most notably for points and axis-parallel rectangles. In fact, it would also be interesting to find a simpler construction that yields *weak ε -nets* of size $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$.

(iii) Last, but not least, there is the problem of showing the existence of small-size ε -nets for the primal range spaces whose duals were considered in section 7, such as those involving planar point sets and locally γ -fat objects, or semiunbounded pseudotrapezoids, with the properties assumed in section 7. (We did achieve this goal for α -fat triangles.)

Appendix. Proof of the Exponential Decay Lemma.

Proof of Lemma 2.1. For a fixed level i , let T denote the collection of all axis-parallel rectangles which are anchored at the entry side of some strip σ_v at that level, and each of their three other sides contains a point of P_v (or extends all the way to the strip boundary or to $\pm\infty$, as appropriate). Let T_t denote the subset of T consisting of all rectangles with weight factor at least t . We have

$$(2) \quad \mathbf{Exp} \{ |\mathbf{CT}_t(R)| \} = \sum_{M \in T_t} \mathbf{Prob} \{ M \in \mathbf{CT}(R) \},$$

$$(3) \quad \mathbf{Exp} \{ |\mathbf{CT}(R')| \} = \sum_{M \in T} \mathbf{Prob} \{ M \in \mathbf{CT}(R') \} \geq \sum_{M \in T_t} \mathbf{Prob} \{ M \in \mathbf{CT}(R') \}.$$

In view of (2) and (3), it suffices to show that, for each $M \in T_t$,

$$\mathbf{Prob} \{ M \in \mathbf{CT}(R) \} = O(2^{-t}) \cdot \mathbf{Prob} \{ M \in \mathbf{CT}(R') \}.$$

Let A_M be the event that $D(M) \subset R$ and $K(M) \cap R = \emptyset$, and let A'_M be the event that $D(M) \subset R'$ and $K(M) \cap R' = \emptyset$. In our set-up, the event A_M is exactly the event $M \in \mathbf{CT}(R)$, and the event A'_M is exactly the event $M \in \mathbf{CT}(R')$. Moreover, putting $\delta := |D(M)| \leq 3$ and $w := |K(M)|$, we have $\mathbf{Prob}\{A_M\} = \pi^\delta(1 - \pi)^w$ and $\mathbf{Prob}\{A'_M\} = (\pi')^\delta(1 - \pi')^w$. Hence

$$\frac{\mathbf{Prob} \{ M \in \mathbf{CT}(R) \}}{\mathbf{Prob} \{ M \in \mathbf{CT}(R') \}} = \frac{\mathbf{Prob} \{ A_M \}}{\mathbf{Prob} \{ A'_M \}} = \frac{\pi^\delta(1 - \pi)^w}{(\pi')^\delta(1 - \pi')^w} = t^\delta \left(\frac{1 - \pi}{1 - \pi'} \right)^w.$$

Substituting $\pi = s/n$, $\pi' = \pi/t$, and $w \geq t \cdot n/s$, the latter expression becomes $O(2^{-t})$, which completes the proof of the lemma. \square

Acknowledgment. The authors wish to thank two anonymous referees for useful comments.

REFERENCES

- [AMS98] P. K. AGARWAL, J. MATOUŠEK, AND O. SCHWARZKOPF, *Computing many faces in arrangements of lines and segments*, SIAM J. Comput., 27 (1998), pp. 491–505.
- [ASS89] P. AGARWAL, M. SHARIR, AND P. SHOR, *Sharp upper and lower bounds for the length of general Davenport Schinzel sequences*, J. Combin. Theory Ser. A, 52 (1989), pp. 228–274.
- [AS92] N. ALON AND J. SPENCER, *The Probabilistic Method*, Wiley Interscience, New York, 1992.
- [BGLR93] M. BELLARE, S. GOLDWASSER, G. LUND, AND A. RUSSELL, *Efficient probabilistically checkable proofs and applications to approximation*, in Proceedings of the 25th Annual ACM Symposium on Theory of Computing, ACM, New York, 1993, pp. 294–304.
- [BEHW89] A. BLUMER, A. EHRENFUCHT, D. HAUSSLER, AND M. WARMUTH, *Classifying learnable geometric concepts with the Vapnik-Chervonenkis dimension*, J. ACM, 36 (1989), pp. 929–965.
- [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 (1998), pp. 485–519.
- [BG95] H. BRÖNNMANN AND M. T. GOODRICH, *Almost optimal set covers in finite VC dimensions*, Discrete Comput. Geom., 14 (1995), pp. 463–479.
- [BL04] H. BRÖNNMANN AND J. LENCHNER, *Fast almost-linear-sized nets for boxes in the plane*, in Proceedings of the 14th Annual Fall Workshop on Computational Geometrics, Boston, MA, 2004, pp. 36–38.

- [CF90] B. CHAZELLE AND J. FRIEDMAN, *A deterministic view of random sampling and its use in geometry*, *Combinatorica*, 10 (1990), pp. 229–249.
- [CCH09] C. CHEKURI, K. CLARKSON, AND S. HAR-PELED, *On the set multi-cover problem in geometric settings*, in *Proceedings of the 25th Annual ACM Symposium on Computational Geometry*, ACM, New York, 2009, pp. 341–350.
- [Cla87] K. L. CLARKSON, *New applications of random sampling in computational geometry*, *Discrete Comput. Geom.*, 2 (1987), pp. 195–222.
- [Cla93] K. L. CLARKSON, *Algorithms for polytope covering and approximation*, in *Proceedings of the 3rd Workshop on Algorithms and Data Structures*, Montreal, Canada, 1993, pp. 246–252.
- [CS89] K. L. CLARKSON AND P. W. SHOR, *Applications of random sampling in computational geometry*, II, *Discrete Comput. Geom.*, 4 (1989), pp. 387–421.
- [CV07] K. L. CLARKSON AND K. VARADARAJAN, *Improved approximation algorithms for geometric set cover*, *Discrete Comput. Geom.*, 37 (2007), pp. 43–58.
- [CLRS01] T. H. CORMEN, C. E. LEISERSON, R. L. RIVEST, AND C. STEIN, *Introduction to Algorithms*, MIT Press, Cambridge, MA, 2001.
- [dB08] M. DE BERG, *Improved bounds on the union complexity of fat objects*, *Discrete Comput. Geom.*, 40 (2008), pp. 127–140.
- [dBCKO08] M. DE BERG, O. CHEONG, M. VAN KREVELD, AND M. OVERMARS, *Computational Geometry: Algorithms and Applications*, 3rd rev. ed., Springer-Verlag, Berlin, 2008.
- [dB10] M. DE BERG, *Better bounds on the union complexity of locally fat objects*, in *Proceedings of the 26th Annual Symposium on Computational Geometry*, Snowbird, UT, 2010, pp. 39–47.
- [EGH*89] H. EDELSBRUNNER, L. GUIBAS, J. HERSHBERGER, J. PACH, R. POLLACK, R. SEIDEL, M. SHARIR, AND J. SNOEYINK, *On arrangements of Jordan arcs with three intersections per pair*, *Discrete Comput. Geom.*, 4 (1989), pp. 523–539.
- [Efr05] A. EFRAT, *The complexity of the union of (α, β) -covered objects*, *SIAM J. Comput.*, 34 (2005), pp. 775–787.
- [EK99] A. EFRAT AND M. KATZ, *On the union of κ -curved objects*, *Comput. Geom. Theory Appl.*, 14 (1999), pp. 241–254.
- [ES00] A. EFRAT AND M. SHARIR, *The complexity of the union of fat objects in the plane*, *Discrete Comput. Geom.*, 23 (2000), pp. 171–189.
- [ERS05] G. EVEN, D. RAWITZ, AND S. SHAHAR, *Hitting sets when the VC-dimension is small*, *Inform. Process. Lett.*, 95 (2005), pp. 358–362.
- [Ezr09] E. EZRA, *Weak ε -Nets for Axis-Parallel Boxes in d -Space*, manuscript, 2009.
- [FG88] T. FEDER AND D. H. GREENE, *Optimal algorithms for approximate clustering*, in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, ACM, New York, 1988, pp. 434–444.
- [Fei98] U. FEIGE, *A threshold of $\ln n$ for approximating set cover*, *J. ACM*, 45 (1998), pp. 634–652.
- [FPT81] R. FOWLER, M. PATERSON, AND S. TANIMOTO, *Optimal packing and covering in the plane are NP-complete*, *Inform. Process. Lett.*, 12 (1981), pp. 133–137.
- [GJ79] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, New York, 1979.
- [HKSS08] S. HAR-PELED, H. KAPLAN, M. SHARIR, AND S. SMORODINSKY, *ε -nets for Halfspaces Revisited*, manuscript.
- [HW87] D. HAUSSLER AND E. WELZL, *ε -nets and simplex range queries*, *Discrete Comput. Geom.*, 2 (1987), pp. 127–151.
- [KRSV08] H. KAPLAN, N. RUBIN, M. SHARIR, AND E. VERBIN, *Efficient colored orthogonal range counting*, *SIAM J. Comput.*, 38 (2008), pp. 982–1011.
- [Kar72] R. M. KARP, *Reducibility among combinatorial problems*, in *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, eds., Plenum Press, New York, 1972, pp. 85–103.
- [KK10] J. KING AND D. KIRKPATRICK, *Improved approximation for guarding simple galleries from the perimeter*, available online at <http://arxiv.org/abs/1001.4231>, 2010.
- [KPW92] J. KOMLÓS, J. PACH, AND G. WOEGINGER, *Almost tight bounds for ε nets*, *Discrete Comput. Geom.*, 7 (1992), pp. 163–173.
- [Mat02] J. MATOÚŠEK, *Lectures on Discrete Geometry*, Springer-Verlag, New York, 2002.
- [Mat92a] J. MATOÚŠEK, *Efficient partition trees*, *Discrete Comput. Geom.*, 8 (1992), pp. 315–334.

- [Mat92b] J. MATOUŠEK, *Reporting points in halfspaces*, *Comput. Geom. Theory Appl.*, 2 (1992), pp. 169–186.
- [Mat95] J. MATOUŠEK, *Approximations and optimal geometric divide-and-conquer*, *J. Comput System Sci.*, 50 (1995), pp. 203–208.
- [MSW90] J. MATOUŠEK, R. SEIDEL, AND E. WELZL, *How to net a lot with little: Small ε -nets for disks and halfspaces*, in *Proceedings of the 6th Annual ACM Symposium on Computational Geometry*, ACM, New York, 1990, pp. 16–22. Revised version available at <http://kam.mff.cuni.cz/~matousek/enets3.ps.gz>.
- [MPSSW94] J. MATOUŠEK, J. PACH, M. SHARIR, S. SIFRONY, AND E. WELZL, *Fat triangles determine linearly many holes*, *SIAM J. Comput.*, 23 (1994), pp. 154–169.
- [NLH84] A. NAAMAD, D. T. LEE, AND W. L. HSU, *On the maximal empty rectangle problem*, *Discrete Appl. Math.*, 8 (1984), pp. 267–277.
- [Niv09] G. NIVASCH, *Improved bounds and new techniques for Davenport-Schinzel sequences and their generalizations*, *J. ACM*, 57 (2010), pp. 1–44.
- [PA95] J. PACH AND P. K. AGARWAL, *Combinatorial Geometry*, Wiley Interscience, New York, 1995.
- [PT02] J. PACH AND G. TARDOS, *On the boundary complexity of the union of fat triangles*, *SIAM J. Comput.*, 31 (2002), pp. 1745–1760.
- [PST95] S. A. PLOTKIN, D. B. SHMOYS, AND É. TARDOS, *Fast approximation algorithms for fractional packing and covering problems*, *Math. Oper. Res.*, 20 (1995), pp. 257–301.
- [PR08] E. PYRGA AND S. RAY, *New existence proofs for ε -nets*, in *Proceedings of the 24th Annual ACM Symposium on Computational Geometry*, ACM, New York, 2008, pp. 199–207.
- [SA95] M. SHARIR AND P. K. AGARWAL, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [Var09] K. VARADARAJAN, *Epsilon nets and union complexity*, in *Proceedings of the 25th Annual ACM Symposium on Computational Geometry*, ACM, New York, 2009, pp. 11–16.
- [You95] N. E. YOUNG, *Randomized rounding without solving the linear program*, in *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, ACM, New York, SIAM, Philadelphia, 1995, pp. 170–178.