Epsilon-Mnets: Hitting Geometric Set Systems with Subsets

Nabil Mustafa, Saurabh Ray

Epsilon-Mnets: Hitting Geometric Set Systems with Subsets*

Nabil H. Mustafa† Saurabh Ray‡

Abstract

The existence of Macbeath regions is a classical theorem in convex geometry [13], with recent applications in discrete and computational geometry. In this paper, we initiate the study of Macbeath regions in a combinatorial setting—and not only for the Lebesgue measure as is the case in the classical theorem—and establish near-optimal bounds for several basic geometric set systems.

1 Introduction

Given a convex body K in $\mathbb{R}^d$ of unit volume, and a parameter $\epsilon > 0$, a classical theorem of Macbeath [13] from convex geometry implies the existence of disjoint convex bodies of K, each of volume $\Theta(\epsilon)$, called *Macbeath regions*, such that any half-space containing at least ϵ-th volume of K completely contains one of these convex bodies. Formally, consider the following theorem (as stated in [6]):

Theorem A (Macbeath regions)**.** *Given a convex body $K \subset \mathbb{R}^d$ of unit volume, and a parameter $0 < \epsilon < 1/(2d)^{2d}$, there exists a set $\mathcal{M}$ of $O\left(\frac{1}{\epsilon^{1-\frac{2}{d+1}}}\right)$ convex objects such that for any half-space h with $\mathrm{vol}(h \cap K) \geq \epsilon$, there exists a $K_i \in \mathcal{M}$ such that $K_i \subset h \cap K$ and*

$$\mathrm{vol}(K_i) \geq \frac{1}{(30d)^d} \cdot \epsilon.$$

Similar partitions of convex bodies was used by Edwald, Larmen and Rogers [9] for cap coverings, which were later further extended by Bárány and Larman [5]. They were also used for lower-bounds on range searching by Brönnimann, Chazelle and Pach [6]. Very recently, Macbeath regions were used in an elegant way by Arya, da Fonseca and Mount [3] for computing near-optimal Hausdorff approximations of polytopes. We refer the reader to Bárány [4] for a survey of these and several other applications of Macbeath regions.

Switching over to discrete and combinatorial geometry, a different structure—ϵ-nets—has been developed over the past three decades as a fundamental and powerful tool in computational geometry. Given a set system $(X, \mathcal{R})$, and a parameter ϵ, an ϵ-net is a set $N \subseteq X$ such that $N \cap R \neq \emptyset$ for all $R \in \mathcal{R}$ with $|R| \geq \epsilon|X|$. A famous theorem of Haussler and Welzl [10] states the existence of ϵ-nets of size $O\big(\frac{d}{\epsilon}\log\frac{d}{\epsilon}\big)$ for $(X, \mathcal{R})$, where d is the VC dimension of $\mathcal{R}$. This bound was later improved in [11] to an optimal bound of $\big(1 + o(1)\big)\big(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\big)$. By now ϵ-nets are an indispensable tool in combinatorics, geometry and algorithms (we refer the reader to the books [20, 16, 7, 17] for a small sampling of their constructions and applications).

---

*A preliminary version of this article appeared in [19].

†Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, Equipe A3SI, ESIEE Paris. E-email: mustafan@esiee.fr. The work of Nabil H. Mustafa in this paper has been supported by the grant ANR SAGA (JCJC-14-CE25-0016-01).

‡Computer Science, New York University, Abu Dhabi. E-email: saurabh.ray@nyu.edu.

The starting point of our work is the observation that the two—ϵ-nets and Macbeath regions—are related. Indeed theorem A implies that for any convex body K in $\mathbb{R}^d$ of volume V, it is possible to pick $O(\frac{1}{\epsilon})$ points in K (in fact, even less) which hit all half-spaces containing an ϵ-th fraction of the volume of K. However, the statement itself is much stronger than that: instead of just points, it states the existence of $O(\frac{1}{\epsilon})$ regions, each of volume $\Theta(\epsilon V)$, so that any half-space containing an ϵ-th fraction of the volume of K contains one of the regions completely. As we will prove in this paper, a strengthening of the ϵ-net statement is true for the counting measure for set systems induced by half-spaces in $\mathbb{R}^3$: given any set P of points in $\mathbb{R}^3$, there exist $O(\frac{1}{\epsilon})$ subsets of P, each of size $\Theta\big(\epsilon|P|\big)$, such that any half-space containing at least $\epsilon \cdot |P|$ points of P contains one of these regions completely. This raises the natural question: of the large number of results known for ϵ-nets for various geometric set systems, which can be optimally strengthened like the case above?

Geometric set systems can be categorized into two frequently studied types. Let $\mathcal{O}$ be a family of geometric objects in $\mathbb{R}^d$—e.g., the family of all half-spaces, all balls and so on. We say that $\mathcal{O}$ has union complexity $\varphi(\cdot)$ if the combinatorial complexity of the union of any r of the regions of $\mathcal{O}$ is at most $r \cdot \varphi(r)$; we refer the reader to the survey [1] for bounds on the union complexity of many geometric objects. Given a set X of points in $\mathbb{R}^d$, we say that $(X, \mathcal{R})$ is a *primal set system* induced by $\mathcal{O}$ if for each $R \in \mathcal{R}$, there exists an object $O \in \mathcal{O}$ such that $R = X \cap O$. On the other hand, given a finite set $\mathcal{S} \subseteq \mathcal{O}$ in $\mathbb{R}^d$, we say that $(\mathcal{S}, \mathcal{R})$ is a *dual set system* induced by $\mathcal{S}$ if for each $R \in \mathcal{R}$, there exists a point $q \in \mathbb{R}^d$ contained in precisely the elements of R, i.e., $R = \big\{O \in \mathcal{S} \mid q \in O\big\}$.

In this paper we initiate a systematic study of the analogues of Macbeath regions—which we name ϵ-Mnets—for some commonly studied primal and dual geometric set-systems.

Definition (ϵ-Mnets)**.** *Given a set system $(X, \mathcal{R})$ and a parameter $\epsilon > 0$, a collection $\mathcal{M} = \{X_1, \ldots, X_t\}$ of subsets of X is an ϵ-Mnet for $\mathcal{R}$ of size t if*

1. *$|X_i| = \Omega\big(\epsilon \cdot |X|\big)$ for each $i = 1, \ldots, t$ and,*
2. *for every $R \in \mathcal{R}$ with $|R| \geq \epsilon \cdot |X|$, there exists an index $j \in \{1, \ldots, t\}$ such that $X_j \subseteq R$.*

Furthermore, for any $\kappa \geq 2$, call $\mathcal{M}$ a $\frac{1}{\kappa}$-heavy ϵ-Mnet if each set in $\mathcal{M}$ has size greater than $\frac{\epsilon|X|}{\kappa}$.

Our Results

Our first result establishes tight bounds for the sizes of ϵ-Mnets for the primal and dual set systems induced by axis-parallel rectangles in the plane. This already provides an example where ϵ-Mnets have larger sizes—by factors polynomial in $\frac{1}{\epsilon}$—than ϵ-nets for the corresponding set systems. The proof of the following statement is in Section 2.

Theorem 1. *Let $\epsilon > 0$, $\kappa \geq 2$ be given parameters.*

(a) ***Dual set system.*** *Given a set $\mathcal{S}$ of axis-parallel rectangles in the plane, there exist $\frac{1}{2^\kappa}$-heavy ϵ-Mnets of size $O\Big(\frac{4^\kappa}{\epsilon^{1+\frac{1}{\kappa}}}\Big)$ for the dual set-system induced by $\mathcal{S}$.*

Furthermore, this is near-optimal: for any integer $n > 0$, there exists a set $\mathcal{S}$ of n axis-parallel rectangles in $\mathbb{R}^2$ such that any $\frac{1}{\kappa}$-heavy ϵ-Mnet for the dual set-system induced by $\mathcal{S}$ has size $\Omega\bigg(\frac{1}{\epsilon^{1+\frac{1}{\kappa-1}}}\bigg)$.

(b) ***Primal set system.*** *Given any set P of points in the plane, there exist ϵ-Mnets of size $O\big(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\big)$ for the primal set-system induced by axis-parallel rectangles on P.*

Furthermore, this is near-optimal: for any integer $n > 0$, there exists a set P of n points in the plane such that any $\frac{1}{\kappa}$-heavy ϵ-Mnet for the primal set-system induced by axis-parallel rectangles on P has size $\Omega\big(\frac{1}{\epsilon}\log_\kappa\frac{1}{\epsilon}\big)$.

Our next result states the existence of small ϵ-Mnets for dual set systems as a function of the union complexity of the objects. Call a set $\mathcal{S}$ of objects in $\mathbb{R}^d$ *well-behaved* if for any subset $\mathcal{S}' \subseteq \mathcal{S}$ and any $Q \subseteq \mathbb{R}^d$, one can decompose the cells in the arrangement of $\mathcal{S}'$ that intersect Q into cells of constant descriptive complexity, where the complexity of this decomposition is proportional to the total number of vertices in the cells that intersect Q; we refer the reader to [8] for more details. The proof of the following statement is in Section 3.

Theorem 2. *Let $\mathcal{R}$ be the dual set system induced by a set of well-behaved regions $\mathcal{S}$ in $\mathbb{R}^d$ with union complexity $\varphi(\cdot)$ and let $\epsilon > 0$ be a given parameter. Then there exists an ϵ-Mnet for $\mathcal{R}$ of size $O\big(\frac{1}{\epsilon}\varphi(\frac{1}{\epsilon})\big)$.*

Interestingly, as $\varphi(m) = \Omega(m)$ for the dual set system induced by axis-parallel rectangles in the plane, Theorem 1 implies that the dependence of $\varphi(\cdot)$ in Theorem 2 cannot be reduced to, for example, $\log \varphi(\cdot)$, as is the case for ϵ-nets.

Our last result is to consider the primal case where the input is a set of points and the set system is defined by containment by geometric objects such as disks, lines, triangles and more generally, k-sided polygons in the plane. The proof of the following statement is in Section 4.

Theorem 3. *Let P be a set of n points, and $\epsilon > 0$ a given parameter. Then one can construct ϵ-Mnets of size:*

(a) $O(\frac{1}{\epsilon^{\lfloor d/2 \rfloor}})$ *for the primal set system induced by half-spaces in $\mathbb{R}^d$, for $d \geq 2$.*

Furthermore, this cannot be improved substantially: for any integers $d \geq 2$ and $n > 0$, there exists a set of n points in $\mathbb{R}^d$ such that any ϵ-Mnet for the primal set system induced by half-spaces has size $\Omega(\frac{1}{\epsilon^{\lfloor \frac{d+1}{3} \rfloor}})$.

(b) $O\big(\frac{1}{\epsilon}\big)$ *for the primal set system induced by disks in the plane.*

(c) $O\big(\frac{1}{\epsilon^3}(\log\frac{1}{\epsilon})^4\big)$ *for the primal set system induced by triangles, and in general k-sided polygons in the plane (the constant in the asymptotic notation depends on k).*

(d) $O\big(\frac{1}{\epsilon^2}(\log\frac{1}{\epsilon})^2\big)$ *for the primal set system induced by lines,* $O\big(\frac{1}{\epsilon^2}(\log\frac{1}{\epsilon})^3\big)$ *for the one induced by cones, and* $O\big(\frac{1}{\epsilon^2}(\log\frac{1}{\epsilon})^4\big)$ *for the one induced by strips in the plane.*

Furthermore, this is near-optimal: for any integer $n > 0$, there exists a set of n points in $\mathbb{R}^2$ such that any ϵ-Mnet for the primal set system induced by lines or cones or strips has size $\Omega\big(\frac{1}{\epsilon^2}\big)$.

(e) $O\big(\frac{1}{\epsilon}\big)$ *for the primal set system induced by axis-parallel rectangles in $\mathbb{R}^2$, all intersecting the y-axis.*

Theorem 3 implies that near-linear bounds for ϵ-Mnets are not possible for even simple primal set-systems such as those induced by lines in the plane. This contrasts sharply with ϵ-net bounds for geometric set systems, which are near-linear for any set system with constant VC dimension.

2 Proof of Theorem 1

The following lemma, of independent interest, gives insight for studying ϵ-Mnets for both the primal and dual set systems induced by axis-parallel rectangles in the plane.

Lemma 2.1. *For any integers $r, d \geq 3$, consider the grid $G = \{0, \cdots, r-1\}^d$ in $\mathbb{R}^d$ consisting of r^d points. Then there exists a bijective mapping $\pi : G \mapsto \mathbb{R}^2$ such that the primal set system on G induced by axis-parallel lines can be realized by the primal set system induced by axis-parallel rectangles in $\mathbb{R}^2$ on the set $\{\pi(p), p \in G\}$.*

Proof. Let $[r]$ represent the set $\{0, \cdots, r-1\}$. For any $i \in \{1, \ldots, d\}$ and integers $a_1, \cdots, a_{i-1}, a_{i+1}, \cdots, a_d \in [r]$, consider the set of points

$$S_i(a_1, \cdots, a_{i-1}, a_{i+1}, \cdots, a_d) = \Big\{(a_1, \cdots, a_{i-1}, t, a_{i+1}, \cdots, a_d) : \ t \in [r]\Big\}.$$

We call such a set a *line in direction* i. There are dr^{d-1} such lines, r^{d-1} in each of the d directions (along the axes) in $\mathbb{R}^d$.

We will show that there exists a mapping $\pi : G \mapsto \mathbb{R}^2$ such that for each line l in any direction, the inclusion-minimal axis-parallel rectangle containing the image, under $\pi(\cdot)$, of the points in l does not contain the image of any other point of G. Here is the mapping $\pi(\cdot)$ that we will use:

$$\pi\big((a_1, \cdots, a_d)\big) = \sum_j a_j \vec{v}_j, \qquad \text{where } \vec{v}_j = (r^j, r^{d+1-j}).$$

For any point $z \in G$, we will interpret $p = \pi(z)$ both as a vector and as a point, as suitable. When treating it as a vector, we will denote it by $\vec{p}$. For any $z' = (a_1, \cdots, a_d) \in G$, let $\vec{V}_{<i}(z')$ denote the vector $\sum_{j<i} a_j \vec{v_j}$ and $\vec{V}_{>i}(z')$ denote the vector $\sum_{j>i} a_j \vec{v_j}$. Thus we can write $\pi(z') = \vec{V}_{<i}(z') + a_i \vec{v}_i + \vec{V}_{>i}(z')$.

Consider any line, say $l = S_i(a_1, \cdots, a_{i-1}, a_{i+1}, \cdots, a_d)$, and let R be the smallest rectangle containing the set of r mapped points of l in the plane, namely the set

$$f(l) = \Big\{\pi\big((a_1, \ldots, a_{i-1}, t, a_{i+1}, \ldots, a_d)\big) : t \in [r]\Big\}.$$

Let $z_l = (a_1, \cdots, a_{i-1}, 0, a_{i+1}, \cdots, a_d)$ and $z_r = (a_1, \cdots, a_{i-1}, r-1, a_{i+1}, \cdots, a_d)$ be the two extreme points lying on l. As all the coordinates except the i-th one are the same for all points lying on l, the mapped point with the maximum x-coordinate is the one that maximizes $t \cdot r^i$, i.e., the point $\pi(z_r)$. Similarly, $\pi(z_r)$ has the maximum y-coordinate, and $\pi(z_l)$ has the minimum x- and y-coordinates. Furthermore, the width of R is defined by the difference in the x-coordinates of $\pi(z_r)$ and $\pi(z_l)$, and so it is precisely $(r-1)r^i$. Likewise, the height of R is $(r-1)r^{d+1-i}$.

It remains to show that for any other point, say $z = (b_1, \cdots, b_d) \in G \setminus l$, $\pi(z)$ does not lie in R. Let $z' = (a_1, \ldots, a_{i-1}, b_i, a_{i+1}, \ldots, a_d) \in G$ be the point lying on the line l with the same i-th coordinate as z. Let $p = \pi(z) = \vec{V}_{<i}(z) + b_i\vec{v_i} + \vec{V}_{>i}(z)$ and $q = \pi(z') = \vec{V}_{<i}(z') + b_i\vec{v_i} + \vec{V}_{>i}(z')$. Then

$$\vec{p} - \vec{q} = \big(\vec{V}_{<i}(z) - \vec{V}_{<i}(z')\big) + \big(\vec{V}_{>i}(z) - \vec{V}_{>i}(z')\big).$$

Since $\vec{p} \neq \vec{q}$, one of the above two summands must be non-zero. Without loss of generality assume that the second summand is non-zero. The other case is similar. As $\vec{V}_{>i}(z) - \vec{V}_{>i}(z') = \sum_{j>i}(b_j - a_j)\vec{v_j}$, it is a non-zero integral combination of the vectors v_j for $j > i$, and so its x-coordinate has magnitude at least r^{i+1}. On the other hand the x-coordinate of $\big(\vec{V}_{<i}(z) - \vec{V}_{<i}(z')\big)$ has magnitude at most $\sum_{1 \leq j < i}(r-1)r^j = r^i - r$. Therefore the difference in the x-coordinates between p and q is at least $r^{i+1} - (r^i - r)$, which is greater than the width of R. Hence, $p \notin R$. When $\big(\vec{V}_{<i}(z) - \vec{V}_{<i}(z')\big) \neq 0$, a similar argument holds for the y-coordinates of p and q, showing that the difference in their y-coordinates is larger than the height of R. □

Case (a): Dual set system.

Lower-bound. We now show that for any integers $\kappa \geq 2$ and $n \geq 0$, there exists a set $\mathcal{R}$ of n axis-parallel rectangles such that any $\frac{1}{\kappa}$-heavy ϵ-Mnet for the dual set system induced by $\mathcal{R}$ has size $\Omega\big(\frac{1}{\epsilon^{1+1/(\kappa-1)}}\big)$. Apply Lemma 2.1 with $d = \kappa$ and $r = \epsilon^{-\frac{1}{d-1}}$. Let G be the grid $[r]^d$ as before. We set $P = \{\pi(p) : p \in G\}$ and let $\mathcal{R}'$ be the set of dr^{d-1} rectangles corresponding to the dr^{d-1} lines in G. Construct the required set $\mathcal{R}$ by replacing each rectangle of $\mathcal{R}'$ with $\frac{\epsilon n}{d}$ copies. Note that $|\mathcal{R}| = \frac{\epsilon n}{d} \cdot dr^{d-1} = n$. Since each of the points in G is contained in d lines (one in each direction), each point of P is contained in d rectangles of $\mathcal{R}'$ and consequently ϵn rectangles of $\mathcal{R}$. Since there is at most one line through two points in G, there are at most $\frac{\epsilon n}{d}$ rectangles of $\mathcal{R}$ that contain any pair of points $p, q \in P$. Since for any $\frac{1}{\kappa}$-heavy ϵ-Mnet $\mathcal{M}$, each $U \in \mathcal{M}$ has size greater than $\frac{\epsilon n}{\kappa}$, it must be that no set in $\mathcal{M}$ can be contained in two sets $\mathcal{R}(p)$ and $\mathcal{R}(q)$ induced by two distinct points p and q in P. Therefore $|\mathcal{M}| \geq |P| = r^d = \epsilon^{-\frac{\kappa}{k-1}} = \frac{1}{\epsilon^{1+\frac{1}{\kappa-1}}}$.

Upper-bound. We now establish an upper-bound for the dual set systems induced by axis-parallel rectangles in the plane.

Construct a hierarchical subdivision on $\mathcal{S}$, as follows. Let $k = \lceil \frac{1}{\epsilon^{1/\kappa}} \rceil$, and for $i = 0, \ldots, \kappa$, set the parameters $n_i = \frac{n}{k^i}$, and $\epsilon_i = \epsilon(\frac{k}{2})^i$. At the 0-th level (here $i = 0$), let $l_1^0, \ldots, l_{k-1}^0$ be a set of $k-1$ vertical lines such that the number of rectangles of $\mathcal{S}$ lying between two consecutive lines—call this region a 'slab'—is at most $\frac{n_0}{k}$. Let $\mathcal{S}_j^0$ be the set of rectangles lying entirely in the j-th slab. For each index $j = 1, \ldots, k-1$, construct a $\frac{\epsilon_0}{4}$-Mnet for all the rectangles of $\mathcal{S}$ intersecting l_j^0. Furthermore, construct an $\epsilon(\frac{k}{2})$-Mnet for the rectangles in $\mathcal{S}_j^0$, for each $j = 1, \ldots, k-1$ in the similar manner as above. The construction continues for κ steps: at the i-level, there are k^i total sub-problems, each sub-problem consists of at most $n_i = \frac{n}{k^i}$ rectangles and with $\epsilon_i = \epsilon(\frac{k}{2})^i$.

At the base case of the recursion, we use a direct $O(\frac{1}{\epsilon_\kappa^2})$-sized construction for the ϵ_κ-Mnet of the k^κ sub-problems at the last κ-level: for the sub-problem of computing a ϵ_κ-Mnet for a set of rectangles $\mathcal{S}'$ where $|\mathcal{S}'| \leq n_\kappa$, construct a set L' of $\frac{8}{\epsilon_\kappa}$ vertical and $\frac{8}{\epsilon_\kappa}$ horizontal lines such that each vertical (resp. horizontal) slab induced by L' contains at most $\frac{\epsilon_\kappa|\mathcal{S}'|}{4}$ vertical (resp. horizontal) boundary edges of the rectangles in $\mathcal{S}'$. For each bounded cell c induced by L', add to $\mathcal{M}$ all the rectangles of $\mathcal{S}'$ completely containing c, if their total number is at least $\frac{\epsilon_k|\mathcal{S}'|}{2}$. Now take any point $q \in \mathbb{R}^2$ lying in at least $\epsilon_\kappa|\mathcal{S}'|$ rectangles of $\mathcal{S}'$ and let c be the cell induced by L' containing q. At least one of the boundary edges of any rectangle R containing q but not containing c must lie in the vertical or horizontal slab induced by L' containing q. Thus there can be only $\frac{\epsilon_\kappa|\mathcal{S}'|}{2}$ such rectangles that contain q but not the cell c. The remaining at least $\frac{\epsilon_\kappa|\mathcal{S}'|}{2}$ rectangles that contain q must then all contain c, and so would form a set in $\mathcal{M}$ of size at least $\frac{\epsilon_\kappa|\mathcal{S}'|}{2}$. Note that the total number of sets added to $\mathcal{M}$ is $O(\frac{1}{\epsilon_k^2})$.

The next two claims conclude the proof by showing that all these Mnets together form an ϵ-Mnet $\mathcal{M}$ for $\mathcal{S}$ of the required size.

Claim 1. *Each set in $\mathcal{M}$ has size $\Theta\Big(\frac{\epsilon n}{2^\kappa}\Big)$. The size of $\mathcal{M}$ is $O\Big(\frac{4^\kappa}{\epsilon^{1+\frac{1}{\kappa}}}\Big)$.*

Proof. At the i-level there are k^i sub-problems, each of size at most $n_i = \frac{n}{k^i}$ with $\epsilon_i = \epsilon(\frac{k}{2})^i$. For each such sub-problem, we partition its set of at most n_i rectangles by $k-1$ lines, and construct a $\frac{\epsilon_i}{4}$-Mnet for the rectangles intersecting these $k-1$ lines. Note that the set of rectangles intersecting any line, and clipped to one side of the line have linear union complexity [1] and by Theorem 2, there exists a $\frac{\epsilon_i}{4}$-Mnet of size $O(\frac{1}{\epsilon_i})$. Hence the total size over all internal sub-problems is:

$$\sum_{i=0}^{\kappa} k^i \cdot (k-1) \cdot O\Big(\frac{1}{\epsilon_i}\Big) \leq \sum_{i=0}^{\kappa} k^{i+1} \cdot O\Big(\frac{2^i}{\epsilon k^i}\Big) = \sum_{i=0}^{\kappa} O\Big(\frac{2^i}{\epsilon^{1+\frac{1}{\kappa}}}\Big) = O\Big(\frac{2^\kappa}{\epsilon^{1+\frac{1}{\kappa}}}\Big).$$

At the last level, after κ steps, we have k^κ sub-problems, each with at most $\frac{n}{k^\kappa}$ rectangles, and $\epsilon_\kappa = \epsilon(\frac{k}{2})^\kappa$. Now use a direct construction which constructs an ϵ-Mnet of size $O(\frac{1}{\epsilon^2})$, to get the total size of Mnet at the last step to be $O(\kappa^k \cdot \frac{1}{\epsilon_k^2}) = O(\frac{4^\kappa}{\epsilon^2 k^\kappa}) = O(\frac{4^\kappa}{\epsilon})$.

At any level i, we construct a ϵ_i-Mnet on a set of at most $\frac{n}{k^i}$ rectangles. So each set in the constructed Mnet has size $\Omega(\epsilon_i \cdot \frac{n}{k^i}) = \Omega(\frac{\epsilon n}{2^i}) = \Omega(\frac{\epsilon n}{2^\kappa})$. □

Claim 2. *For each point $q \in \mathbb{R}^2$ lying in at least ϵn rectangles of $\mathcal{S}$, there exists a set $U \in \mathcal{M}$ such that q lies in all the rectangles of U.*

Proof. Take a point q lying in at least ϵn rectangles of $\mathcal{S}$. At the 0-th level, say q lies in the vertical slab defined by lines l_j^0 and l_{j+1}^0. If q is contained in at least $\frac{\epsilon n}{4}$ rectangles intersected by either l_j^0 or l_{j+1}^0, say l_j^0, then it is contained in at least $\frac{\epsilon n}{4}$ rectangles out of a total of at most n rectangles intersected by l_j^0. So the $\frac{\epsilon}{4}$-Mnet for l_j^0 will have a set U such that each rectangle in U contain q. Otherwise q is contained in at least $\frac{\epsilon n}{2} = \epsilon(\frac{k}{2})(\frac{n}{k}) = \epsilon_1 n_1$ rectangles of the set $\mathcal{S}_j^0$ of size at most $n_1 = \frac{n_0}{k}$, and we proceed to this sub-problem.

In general, at the i-level, each sub-problem has at most $n_i = \frac{n}{k^i}$ rectangles, with $\epsilon_i = \epsilon(\frac{k}{2})^i$. Then either q is contained in at least $\frac{\epsilon_i n_i}{4}$ rectangles intersecting one of the lines, and so will contain a set from the $\frac{\epsilon_i}{4}$-Mnet constructed for each of the $k-1$ vertical lines. Or q is contained in at least $\frac{\epsilon_i n_i}{2}$ rectangles out of a total of at most $n_{i+1} = \frac{n_i}{k}$ rectangles lying in one of the slabs defined by the $k-1$ vertical lines. But as

$$\frac{\epsilon_i n_i}{2} = \frac{\epsilon}{2} \cdot \left(\frac{k}{2}\right)^i \cdot \frac{n}{k^i} = \epsilon \left(\frac{k}{2}\right)^{i+1} \frac{n}{k^{i+1}} = \epsilon_{i+1} n_{i+1},$$

q will be covered inductively by the ϵ_{i+1}-Mnet constructed for the n_{i+1} rectangles in one of the resulting sub-problems at level $i+1$. □

Case (b): Primal set system.

Lower-bound. We now show that for any integers $\kappa \geq 2$ and $n \geq 0$, there exists a set P of n points in $\mathbb{R}^2$ such that any $\frac{1}{\kappa}$-heavy ϵ-Mnet of P for the primal set system induced by axis-parallel rectangles in the plane has size $\Omega(\frac{1}{\epsilon} \log_\kappa \frac{1}{\epsilon})$. Apply Lemma 2.1 with $r = \kappa$, and with parameter d set with $r^{d-1} = \frac{1}{\epsilon}$. According to the lemma, there is a mapping π from the grid $G = \{0, \cdots, r-1\}^d$ to the plane so that for each subset $S \subset G$ of the grid obtained by intersecting G with an axis-parallel line, there exists an axis-parallel rectangle R in the plane such that $R \cap \pi(G) = \pi(S)$; i.e., R contains exactly the mapped points of S. There are $dr^{d-1} = \Theta(\frac{1}{\epsilon} \log_\kappa \frac{1}{\epsilon})$ such subsets and let $\mathcal{R}$ be the set of axis-parallel rectangles corresponding to these. Let P be the set of points obtained by replacing each point $p \in \pi(G)$ with $\frac{\epsilon n}{r}$ copies of p (note that P is not a multi-set; think of each copy of the same point in $\pi(G)$ as a distinct point). The number of points in P is $r^{d-1} \cdot \frac{\epsilon n}{r} = n$. Each rectangle in $\mathcal{R}$ contains $r \cdot \frac{\epsilon n}{r} = \epsilon n$ points of P. Also, any pair of rectangles in $\mathcal{R}$ share at most $\frac{\epsilon n}{r} = \epsilon \frac{n}{k}$ points of P. Thus no two rectangles in $\mathcal{R}$ may share the same set $U \in \mathcal{M}$ of a $\frac{1}{k}$-heavy ϵ-Mnet $\mathcal{M}$. Since each of them must contain some $U \in \mathcal{M}$, we have $|\mathcal{M}| \geq |\mathcal{R}|$ and the result follows.

Upper-bound. We now present a matching upper-bound for the primal set system induced by axis-parallel rectangles in the plane.

Assume $P = \{p_1, \ldots, p_n\}$ are labeled in the order of increasing x-coordinates. Given P, construct a balanced binary subdivision of P with vertical lines: divide P by a vertical line into two equal-sized subsets P_0^1, P_1^1, and then recursively divide each of these sets into two equal-sized subsets and so on for $\log \frac{1}{\epsilon}$ levels. At the i^{th} level of recursion, there are 2^i sets of size $\frac{n}{2^i}$.

Let P_j^i denote the j-th subset of P at level i, i.e.,

$$\text{For } 1 \leq i \leq \log \frac{1}{\epsilon}, \quad 0 \leq j < 2^i, \quad P_j^i = \left\{ p_{j\frac{n}{2^i}+1}, \ldots, p_{(j+1)\frac{n}{2^i}} \right\}.$$

For each set P_j^i, and for each of its two bounding lines, say lines l_0 and l_1, construct a $2^{i-1}\epsilon$-Mnet for the following primal set-system: the base set is P_j^i, and given a line $l \in \{l_0, l_1\}$, the sets are induced by axis-parallel rectangles intersecting the line l. Note that all points of P_j^i lie on the same side of l. Let $\mathcal{M}$ be the union of all these Mnets. Crucially, the primal set system induced by the set of axis-parallel rectangles on the same side of l admits an ϵ-Mnet of size $O(\frac{1}{\epsilon})$ by Theorem 3 (e).

We now prove that $\mathcal{M}$ is an ϵ-Mnet of P, of size $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$.

Claim 3. *Each set in $\mathcal{M}$ has size $\Theta(\epsilon n)$, and size of $\mathcal{M}$ is $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$.*

Proof. The set P^i_j has size $\frac{n}{2^i}$, and so each set in a $(2^{i-1}\epsilon)$-Mnet of P^i_j has size $\Omega(2^{i-1}\epsilon\cdot\frac{n}{2^i}) = \Omega(\epsilon n)$. Note that each $2^{i-1}\epsilon$-Mnet has size $O(\frac{1}{2^i\epsilon})$, there are 2^i sets P^i_j at level i, and a total of $\log\frac{1}{\epsilon}$ levels. Hence the size of $\mathcal{M}$ is $O(\frac{1}{2^i\epsilon}\cdot 2^i\cdot\log\frac{1}{\epsilon}) = O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$. □

Claim 4. *Each axis-parallel rectangle containing at least ϵn points of P contains a set of $\mathcal{M}$.*

Proof. Let R be an axis-parallel rectangle containing at least ϵn points of P. Let i be the smallest index such that R intersects exactly one vertical line separating two sets P^i_j and P^i_{j+1} at level i. Say R intersects the line l separating P^i_j and P^i_{j+1}. Then R must contain at least $\frac{\epsilon n}{2}$ points from either P^i_j or P^i_{j+1}, say P^i_j. Let R' be the part of R on the side of l towards P^i_j. Thus R' must contain at least one set of the $2^{i-1}\epsilon$-Mnet for P^i_j, as

$$|R\cap P^i_j| = |R'\cap P^i_j| \geq \frac{\epsilon n}{2} = 2^{i-1}\epsilon\cdot\frac{n}{2^i} = 2^{i-1}\epsilon\cdot|P^i_j|.$$

□

3 Proof of Theorem 2

Given the input set $\mathcal{S}$ of regions in $\mathbb{R}^d$, define the *depth* of any point $q\in\mathbb{R}^d$ with respect to $\mathcal{S}$ to be the number of regions of $\mathcal{S}$ containing q. The key tool used in the proof are *shallow cuttings*:

Theorem B ([15, 8])**.** *Given a set $\mathcal{S}$ of n well-behaved regions in $\mathbb{R}^d$ with union complexity $\varphi(\cdot)$ and two parameters $r, l > 0$, there exists a partition of $\mathbb{R}^d$ into a set Ξ of interior-disjoint cells (of constant description complexity) such that*

1. *each cell of Ξ is intersected by the boundary of at most $\frac{n}{r}$ regions of $\mathcal{S}$, and*
2. *the number of cells in Ξ that contain points of depth less than l (with respect to $\mathcal{S}$) is $O\Big(\big(\frac{rl}{n}+1\big)^d\cdot\frac{n}{l}\cdot\varphi\big(\frac{n}{l}\big)\Big)$.*

Such a partition Ξ is called a $(\frac{1}{r}, l)$-shallow cutting of $\mathcal{S}$.

We will construct the required ϵ-Mnet $\mathcal{M}$ as a union of $\log\frac{1}{\epsilon}$ collections $\mathcal{M}_i$, for $i = 0,\ldots,\log\frac{1}{\epsilon}$. For a fixed index i, construct the sets in $\mathcal{M}_i$ by setting $l_i = 2^{i+1}\epsilon n$, $r_i = \frac{1}{2^{i-1}\epsilon}$, and construct a $(\frac{1}{r_i}, l_i)$-shallow cutting, denoted by Ξ_i, for $\mathcal{S}$. Call a cell $\Delta\in\Xi_i$ *shallow* if it contains points of depth less than l_i. For each $\Delta\in\Xi_i$, let $r(\Delta)$ be the set of regions in $\mathcal{S}$ that completely contain Δ; i.e., $S\in r(\Delta)$ if and only if $\Delta\subset S$. Now, for all shallow cells Δ with $r(\Delta)\geq\frac{\epsilon n}{2}$, add $r(\Delta)$ to $\mathcal{M}_i$.

We can trivially upper-bound $|\mathcal{M}_i|$ by the number of shallow cells of Ξ_i, i.e., cells containing a point of depth less than $l_i = 2^{i+1}\epsilon n$. Thus using Theorem B, we get

$$|\mathcal{M}_i| = O\Bigg(\Big(\frac{r_i\cdot 2^{i+1}\epsilon n}{n}+1\Big)^d\cdot\frac{n}{2^{i+1}\epsilon n}\cdot\varphi\Big(\frac{n}{2^{i+1}\epsilon n}\Big)\Bigg) = O\Big(4^d\cdot\frac{1}{2^i\epsilon}\cdot\varphi\big(\frac{1}{2^i\epsilon}\big)\Big).$$

First we bound the size of $\mathcal{M} = \bigcup_i \mathcal{M}_i$:

$$|\mathcal{M}| \leq \sum_{i=0}^{\log \frac{1}{\epsilon}} |\mathcal{M}_i| = \sum_{i=0}^{\log \frac{1}{\epsilon}} O\Big(4^d \cdot \frac{1}{2^i \epsilon} \cdot \varphi(\frac{1}{2^i \epsilon})\Big) = O\Big(4^d \cdot \frac{1}{\epsilon} \cdot \varphi(\frac{1}{\epsilon})\Big) \sum_{i=0}^{\log \frac{1}{\epsilon}} \frac{1}{2^i} = O\Big(4^d \cdot \frac{1}{\epsilon} \cdot \varphi(\frac{1}{\epsilon})\Big).$$

To see that sets in $\mathcal{M}$ form the required ϵ-Mnet, let $p \in \mathbb{R}^d$ be any point contained in t regions of $\mathcal{S}$, where $t \geq \epsilon n$. Let i be the index such that $2^i \epsilon n \leq t < 2^{i+1} \epsilon n$. Let Δ_p be the shallow cell in the $(\frac{1}{r_i}, l_i)$-shallow cutting that contains p. Recall that the $(\frac{1}{r_i}, l_i)$-shallow cutting Ξ_i partitions $\mathbb{R}^d$ into a set of cells such that each cell intersects the boundary of at most $\frac{n}{r_i} = 2^{i-1}\epsilon n$ objects in $\mathcal{S}$. Thus, of all the $t \geq 2^i \epsilon n$ regions containing p, the boundary of at most $2^{i-1}\epsilon n$ regions can intersect Δ_p. The remaining at least $2^{i-1}\epsilon n \geq \frac{\epsilon n}{2}$ regions of $\mathcal{S}$ containing p must then completely contain Δ_p, and so are in the set $r(\Delta_p)$. Thus the set $r(\Delta_p)$ is added to $\mathcal{M}_i$, and we have $\big|r_i(\Delta)\big| \geq 2^{i-1}\epsilon n \geq \frac{\epsilon n}{2}$.

4 Proof of Theorem 3

(a). First we establish the upper-bound on the sizes of ϵ-Mnets for the primal set system induced by half-spaces in $\mathbb{R}^d$. For a point $p \in P$, let H_p be its dual hyperplane, and let $\mathcal{H} = \{H_p \mid p \in P\}$. Let $\mathcal{H}^+$ (resp. $\mathcal{H}^-$) be a set of upperward-facing (resp. downward-facing) half-spaces defined by $\mathcal{H}$. Apply Theorem 2 to the dual set system induced by $\mathcal{H}^+$ (resp. H^-) to get an ϵ-Mnet $\mathcal{M}^+$ (resp. $\mathcal{M}^-$), and let $\mathcal{M}$ be the corresponding collection of sets for P corresponding to both $\mathcal{M}^+$ and $\mathcal{M}^-$. As $\mathcal{M}^+$ (resp. $\mathcal{M}^-$) is an ϵ-Mnet for $\mathcal{H}^+$ (resp. $\mathcal{H}^-$), for any point $q \in \mathbb{R}^d$ contained in at least ϵn half-spaces in $\mathcal{H}^+$ (resp. $\mathcal{H}^-$), there exists a set in $\mathcal{M}^+$ (resp. $\mathcal{M}^-$) of size $\Omega(\epsilon n)$, such that each half-space in this set contains q. Switching to the primal viewpoint, any upward-facing (resp. downward-facing) half-space H_q containing at least ϵn points of P, corresponds in the dual to a point q that is contained in at least ϵn downward-facing (resp. upward-facing) half-spaces in $\mathcal{H}^+$ (resp. $\mathcal{H}^-$). As $\mathcal{M}^+$ (resp. M^-) is an ϵ-Mnet for $\mathcal{H}^+$ (resp. $\mathcal{H}^-$), it follows that $\mathcal{M}$ is an ϵ-Mnet for the primal set system induced by half-spaces. To bound the size of $\mathcal{M}$ obtained from Theorem 2, it suffices to note that for half-spaces, $r\varphi(r) = O(r^{\lfloor d/2 \rfloor})$ [1].

For the lower-bound for ϵ-Mnets for the primal set system induced by half-spaces in $\mathbb{R}^d$, we first prove the following more general theorem.

Theorem 4. *Given a real parameter $\epsilon > 0$, integer $n > 1$ and two constants δ and k, there exists a set P of n points in the plane, and a set $\mathcal{D}$ of $\Omega(\frac{1}{\epsilon^{\delta+1}})$ curves, each of degree at most δ, such that $a)$ each curve contains ϵn points of P and $b)$ no two curves in $\mathcal{D}$ have more than $\frac{\epsilon n}{k}$ points of P in common. In particular, any $\frac{1}{k}$-heavy ϵ-Mnet for the primal set system on P induced by curves of degree at most δ has size $\Omega(\frac{1}{\epsilon^{\delta+1}})$ (the constants in the asymptotic notation depend on k and δ).*

Proof. Denote by G the set of $\frac{\delta k}{\epsilon}$ grid points in $\{0, \ldots, \delta k - 1\} \times \{0, \ldots, \lceil \frac{1}{\epsilon} \rceil - 1\}$. The set of curves in $\mathcal{D}$ will be all univariate functions in x of the form

$$y = \sum_{i=0}^{\delta} a_i \cdot x^i, \quad \text{where each } a_i \in \Big\{0, 1, \ldots, \lceil \frac{1}{\epsilon(\delta+1)(\delta k)^i} \rceil - 1\Big\}.$$

Clearly we have

$$|\mathcal{D}| = \prod_{i=0}^{\delta} \frac{1}{\epsilon(\delta+1)(\delta k)^i} = \Omega\Big(\frac{1}{\epsilon^{\delta+1}(\delta k)^{\Theta(\delta^2)}}\Big) = \Omega\big(\frac{1}{\epsilon^{\delta+1}}\big).$$

Since for each value of $x \in \{0, \ldots, \delta k - 1\}$, the corresponding value of y for each of the curves in $\mathcal{D}$ lies in $\{0, \ldots, \lceil \frac{1}{\epsilon} \rceil - 1\}$, each of the curves of $\mathcal{D}$ contain precisely δk points of G. Furthermore, as these curves have degree at most δ, no two intersect in more than δ points of G.

Let P be the set of n points obtained by replacing each point of G with $\frac{\epsilon n}{\delta k}$ copies to get a set of n points in the plane. Now each curve in $\mathcal{D}$ contains $\delta k \cdot \frac{\epsilon n}{\delta k} = \epsilon n$ points of P and every pair of curves have less than $d \cdot \frac{\epsilon n}{\delta k} = \frac{\epsilon n}{k}$ points of P in common.

Finally observe that any $\frac{1}{k}$-heavy ϵ-Mnet $\mathcal{M}$ for the primal set system on P induced by $\mathcal{D}$ must consist of at least $|\mathcal{D}|$ sets: each curve $D \in \mathcal{D}$ must completely contain a set $R \in \mathcal{M}$ of size at least $\frac{\epsilon n}{k}$, and furthermore R cannot be contained in any other curve $D' \in \mathcal{D}$, as any two curves of $\mathcal{D}$ have less than $\frac{\epsilon n}{k}$ points of P in common. ☐

Now we show the desired lower-bound for ϵ-Mnets for the primal set system induced by half-spaces in $\mathbb{R}^d$.

Corollary 4.1. *For any $\epsilon > 0$ and integers n and d, there exists a set P of n points in $\mathbb{R}^d$ such that any ϵ-Mnet for the primal set system on P induced by half-spaces has size $\Omega\left(\frac{1}{\epsilon^{\lfloor \frac{d+1}{3} \rfloor}}\right)$.*

Proof. First assume that $\frac{d-2}{3}$ is an integer, and apply Theorem 4 with $\delta = \frac{d-2}{3}$ and $k = 2$ to get a set P of n points in $\mathbb{R}^2$ and a set $\mathcal{D}$ of curves such that any ϵ-Mnet for the primal set system induced by $\mathcal{D}$ on P has size $\Omega(\frac{1}{\epsilon^{\delta+1}})$. We now use Veronese maps [17] to map the incidences between points and curves in $\mathcal{D}$ to incidences between points and half-spaces in $\mathbb{R}^d$. More precisely, consider the map:

$$\pi : p = (p_x, p_y) \in \mathbb{R}^2 \longrightarrow (x, x^2, \ldots, x^{2\delta}, y, yx, \ldots, yx^{\delta}, y^2) \in \mathbb{R}^d.$$

We claim that for any curve $D \in \mathcal{D}$, say defined by the equation $y = \sum_{i=0}^{\delta} a_i \cdot x^i$, there exists a half-space H_D in $\mathbb{R}^d$ such that the set of points of P contained in D is precisely the set of points of $\pi(P)$ contained in H_D. The required half-space can be constructed as follows:

$$\begin{aligned} p \in D \quad \text{if and only if} \quad & \left(y - \sum_{i=0}^{\delta} a_i \cdot x^i\right) = 0 \\ & \left(y - \sum_{i=0}^{\delta} a_i \cdot x^i\right)^2 \leq 0 \\ & \left(a'_1 x + a'_2 x^2 + \cdots + a'_{2\delta} x^{2\delta}\right) + \Big(-2y \cdot \left(a_0 x^0 + \cdots + a_\delta x^\delta\right)\Big) + y^2 \leq a'_0 \end{aligned}$$

for constants $a'_0, \ldots, a'_{2\delta}$ depending on $a_0, \ldots, a_\delta$. Labeling the coordinates in $\mathbb{R}^{3\delta+2}$ with $x_1, \ldots, x_{3\delta+2}$, the required half-space H_D is then

$$H_D : \quad a'_1 \cdot x_1 + \cdots + a'_{2\delta} \cdot x_{2\delta} + (-2a_0) \cdot x_{2\delta+1} + \cdots + (-2a_\delta) x_{3\delta+1} + x_{3\delta+2} \leq a'_0,$$

containing precisely the points that lie on the curve $D \in \mathcal{D}$. This now implies a lower-bound of $\Omega(\frac{1}{\epsilon^{\delta+1}}) = \Omega(\frac{1}{\epsilon^{(d+1)/3}})$ for the ϵ-Mnet for the primal set system induced by half-spaces in $\mathbb{R}^d$. Finally, the lower-bound follows for any value of d by applying the bound for the largest $d' \leq d$ with integer value of $\frac{d'-2}{3}$. ☐

(b). By Veronese maps, points P and disks D can be lifted to half-spaces H in $\mathbb{R}^3$ such that each point is lifted to a point in $\mathbb{R}^3$ and each disk is lifted to a half-space in $\mathbb{R}^3$ in such a way that their incidences are preserved. Now the required upper-bound follows from applying the bound in part **(a)** for half-spaces in $\mathbb{R}^3$ to the lifted point set of P.

(c). As a k-sided polygon can be partitioned into k triangles, one of which must contain at least $\frac{\epsilon n}{k}$ points,

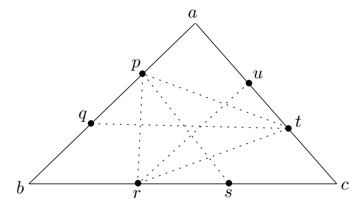

an $\frac{\epsilon}{k}$-Mnet with respect to triangles is an ϵ-Mnet with respect to k-sided polygons. Thus from now on we restrict ourselves to the primal set system induced by triangles in the plane.

Consider any triangle T in the plane that contains ϵn points of P. By moving the sides of the triangle we can ensure that each side of T contains at least two points of P and this can be done in such a way that no point outside T enters the interior of P. Some points in the interior of T may have moved to its boundary and some point outside T may also have moved to the boundary. Since at most 6 points may be on the boundary of T, due to P being in general position, the interior of T still contains at least $\frac{\epsilon n}{2}$ points, assuming $\epsilon n \geq 12$ (observe that for $\epsilon n < 12$, the collection of singletons of P is an ϵ-Mnet of size $O(\frac{1}{\epsilon})$). Thus we can further restrict ourselves to the interior of triangles each of whose sides contain at least two points. The figure above shows a triangle with each side containing two points of P. The points q and r could be identical, they could both be equal at the corner b of the triangle. Similarly s and t could be at c and u and p could be at a. Observe that the triangles aqt, bsp, cur and prt cover the triangle T and therefore one of them must contain at least $\frac{\epsilon n}{4}$ points of P. Each of these triangles are of the following type: at least two of the corners are in P and all sides contain at least two points of P. We call such triangles *anchored* triangles. Thus we can again restrict ourselves to the problem of anchored triangles in the plane containing ϵn points of P.

Let $\mathcal{O}$ be the set of all anchored triangles for P. Let $\mathcal{O}' = \{\Delta_1, \ldots, \Delta_t\}$ be a maximal set of t triangles from $\mathcal{O}$ such that $|\Delta_i \cap P| = \epsilon n$ and $|\Delta_i \cap \Delta_j \cap P| \leq \frac{\epsilon n}{2}$.

Lemma 4.1. $|\mathcal{O}'| \leq 2 \cdot f_{\mathcal{O}}\big(\frac{c}{\epsilon} \cdot \log \frac{1}{\epsilon}, 2c \log \frac{1}{\epsilon}\big)$*, where* $f_{\mathcal{O}}(m, l)$ *is the maximum number of subsets of size at most* l *in the primal set system induced by objects in* $\mathcal{O}$ *on any subset of* m *points of* P*, and* c *is some fixed constant.*

Proof. Pick each point of P independently at random with probability $p = \frac{c}{2\epsilon n} \cdot \log \frac{1}{\epsilon}$ to get a random sample S.

First, observe that with probability greater than $\frac{1}{2}$, the sets $\Delta_i \cap S$, $i = 1 \ldots t$, are distinct and $|S| \leq \frac{c}{\epsilon} \cdot \log \frac{1}{\epsilon}$: consider the range space $(P, \mathcal{R}')$, where $\mathcal{R}' = \big\{(\Delta_i \setminus \Delta_j) \cap P \mid \forall 1 \leq i < j \leq t\big\}$. From the definition of $\mathcal{O}'$, each set in $\mathcal{R}'$ has size at least $\epsilon n - \frac{\epsilon n}{2} = \Theta(\epsilon n)$. We now use the fact that ranges induced by polygons with k sides have VC dimension at most $2k + 1$ [17]; it is easy to see that $\mathcal{R}'$ is a subset of the ranges induced by polygons (or union of polygons) with at most 9 sides, and so the VC dimension of $\mathcal{R}'$ is at most 19. Then by the Haussler-Welzl theorem [10], for $c > 19 \cdot 4$, with probability greater than $\frac{3}{4}$, S is an ϵ-net for $(P, \mathcal{R}')$. Now observe that if $\Delta_i \cap S = \Delta_j \cap S$, then the set $(\Delta_i \setminus \Delta_j) \cap S$ is empty, a contradiction to the fact that S is an ϵ-net for $\mathcal{R}'$. From standard concentration estimates from Chernoff bounds, it follows that $|S| \geq \frac{c}{\epsilon} \cdot \log \frac{1}{\epsilon}$ with probability less than $\frac{1}{4}$.

For each $\Delta_i \in \mathcal{O}'$, let X_i be the random variable which is 1 if $|\Delta_i \cap S| \geq 2c \log \frac{1}{\epsilon}$, and 0 otherwise. For a fixed i, by linearity of expectation, we have $E[|\Delta_i \cap S|] = \frac{c}{2} \cdot \log \frac{1}{\epsilon}$. By Markov's inequality applied to each X_i,

$$\Pr[X_i = 1] = \Pr[|\Delta_i \cap S| \geq 2c \cdot \log \frac{1}{\epsilon}] = \Pr\big[|\Delta_i \cap S| \geq 4 \cdot \mathbb{E}[|\Delta_i \cap S|]\big] \leq \frac{1}{4}.$$

Hence $\mathbb{E}[Y] = \mathbb{E}[\sum X_i] \leq \frac{t}{4}$, and by Markov's inequality applied to Y, we get that $\Pr\big[\sum X_i \geq \frac{t}{2}\big] \leq \frac{1}{2}$.

We can conclude that there exists a subset S of size $\frac{c}{\epsilon} \log \frac{1}{\epsilon}$ such that $\Delta_i \cap S$ are distinct for all objects in $\mathcal{O}'$, and for at least $\frac{|\mathcal{O}'|}{2}$ of the objects in $\mathcal{O}'$, we have $|\Delta_i \cap S| \leq 2c \log \frac{1}{\epsilon}$. Therefore we can get the required bound on the size of $\mathcal{O}'$:

$$\frac{|\mathcal{O}'|}{2} \leq f\big(|S|, l\big) = f\big(\frac{c}{\epsilon} \log \frac{1}{\epsilon}, 2c \log \frac{1}{\epsilon}\big).$$

□

Remark: After the appearance of the conference version of this paper, the statement of Lemma 4.1 has been formalized as the *shallow packing lemma*. We refer the reader to [18] for details and recent history.

We will need the following theorem from [14].

Theorem C (Simplicial partition theorem). *Given a set P of n points in $\mathbb{R}^d$, and an integer parameter $t > 0$, there exists a partition of P into t sets, each of size $\Theta(\frac{n}{t})$, such that any hyperplane intersects the convex-hull of at most $O(t^{1-1/d})$ sets of the partition.*

Take this set $\mathcal{O}'$ of maximal objects, each containing ϵn points of P, and every pair of objects in $\mathcal{O}'$ intersecting in less than $\frac{\epsilon n}{2}$ points. For each object $\Delta_i \in \mathcal{O}'$, do the following: apply the simplicial partition theorem to $\Delta_i \cap P$ with the parameter t, set to a large enough constant, to get a partition of $\Delta_i \cap P$ into t sets of size $\Theta(\frac{|\Delta_i \cap P|}{t})$. Add each of these $t = O(1)$ sets to the ϵ-Mnet $\mathcal{M}$ for P.

Claim 5. *$\mathcal{M}$ is an ϵ-Mnet for the primal set-system induced by $\mathcal{O}$, of size $O\Big(f_{\mathcal{O}}\big(\frac{c}{\epsilon} \cdot \log \frac{1}{\epsilon}, 2c \log \frac{1}{\epsilon}\big)\Big)$.*

Proof. First note that each set added to $\mathcal{M}$ had size $\Theta(\frac{|\Delta_i \cap P|}{t}) = \Theta(\epsilon n)$, and the number of such sets is $O(|\mathcal{O}'| \cdot t) = O(|\mathcal{O}'|)$. It remains to show that any object containing ϵn points of P contains one set of $\mathcal{M}$. Take any triangle Δ containing ϵn points of P (any triangle containing greater than ϵn points can always be shrunk to a triangle containing fewer points). By the maximality of $\mathcal{O}'$, there exists $\Delta_i \in \mathcal{O}'$ such that $|\Delta \cap \Delta_i| \geq \frac{\epsilon n}{2}$. Furthermore, of all the sets in the simplicial partition of Δ_i, each edge of $\partial\Delta$ can intersect only $O(\sqrt{t})$ sets; so in total the three bounding segments of Δ can intersect at most $O(3\sqrt{t})$ sets. Each of these sets has $O(\frac{|\Delta_i \cap P|}{t})$ points. So these sets can contribute at most $O(3\sqrt{t} \cdot \frac{|\Delta_i \cap P|}{t})$ points of Δ_i to Δ. Setting t to be a large-enough constant (say, $t = 38$), this is less than $\frac{\epsilon n}{2}$. Therefore Δ must contain a point in Δ_i which lies in a partition for Δ_i not intersecting $\partial\Delta$, i.e., the partition lies completely inside Δ. ☐

Finally, when $\mathcal{O}$ is a set of anchored triangles in the plane, a routine application of the Clarkson-Shor method [17] implies that $f_{\mathcal{O}}(n, l) = O(n^3 \cdot l)$. Then Lemma 5 implies the existence of ϵ-Mnets for the primal set system induced by $\mathcal{O}$ of size $O\big((\frac{c}{\epsilon} \log \frac{1}{\epsilon})^3 \cdot 2c \log \frac{1}{\epsilon}\big) = O\big(\frac{1}{\epsilon^3}(\log \frac{1}{\epsilon})^4\big)$.

(d). The upper-bounds for the primal set systems induced by lines, strips, cones in the plane again follow from Lemma 5. The function $f(n, l)$ correspondingly denotes the number of subsets of size l induced by the objects of the appropriate type (lines, strips, cones). For lines, $f(n, l) = O(n^2)$ implies the existence of ϵ-Mnets of size $O(\frac{1}{\epsilon^2}(\log \frac{1}{\epsilon})^2)$; for strips $f(n, l) = O(n^2 \cdot l)$ implies the existence of ϵ-Mnets of size $O(\frac{1}{\epsilon^2}(\log \frac{1}{\epsilon})^3)$; and for cones, $f(n, l) = O(n^2 \cdot l^2)$ implies the existence of ϵ-Mnets of size $O(\frac{1}{\epsilon^2}(\log \frac{1}{\epsilon})^4)$.

The lower-bound for the primal set system induced by lines, strips and cones in the plane follows from Theorem 4 by setting $\delta = 1$:

Corollary 4.2. *For any $\epsilon > 0$ and integer n, there exists a set P of n points in the plane such that any ϵ-Mnet for the primal set system on P induced by lines must have size $\Omega(\frac{1}{\epsilon^2})$.*

As the set system induced by lines is a special case for the ones induced by strips and cones, this implies the same lower-bound for the primal set system induced by strips and cones in the plane.

(e). As each rectangle contains ϵn points of P and intersects the y-axis, for each rectangle R, take the portion of the rectangle on the side of the y-axis that contains at least $\frac{\epsilon n}{2}$ points. We can construct $\frac{\epsilon}{2}$-Mnets for the two sides of the y-axis separately and return the union of the two Mnets. Now for the primal set system induced by axis-parallel rectangles with one vertical edge lying on the y-axis, we have $f(n, l) = O(n)$ [22]. Now Lemma 5 implies that one can construct $\frac{\epsilon}{2}$-Mnets of size $O(\frac{1}{\epsilon})$.

5 Conclusion and future work

We conclude our study by observing that the above series of results— with proofs that use different techniques— indicate an intriguing relation between the sizes of ϵ-nets and the sizes of ϵ-Mnets. In all cases, they obey the following pattern: if there exist ϵ-nets of size $O\big(\frac{1}{\epsilon}f(\frac{1}{\epsilon})\big)$ for some primal or dual set system, then the size of ϵ-Mnets for the same set system is $O(\frac{1}{\epsilon}c^{f(\frac{1}{\epsilon})})$, where c is some constant. For example, for all spaces known to have linear-sized ϵ-nets (which is optimal), our proofs establish the existence of linear-sized ϵ-Mnets (which is optimal). For the primal set system induced by axis-parallel rectangles in the plane, ϵ-nets have size $O(\frac{1}{\epsilon}\log\log\frac{1}{\epsilon})$ (shown to be optimal) [2, 21]; our results show the existence of ϵ-Mnets of size $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ (which we show to be optimal). For the primal set system induced by half-spaces in $\mathbb{R}^d$, ϵ-nets have size $O(\frac{d}{\epsilon}\log\frac{1}{\epsilon})$ (shown to be optimal [12]); our results establish the existence of ϵ-Mnets for this set system of size $O(\frac{1}{\epsilon^{(d+1)/3}})$. Similarly, for the remaining set systems for which there exist ϵ-nets of size $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$, we show the existence of ϵ-Mnets of size $O(\frac{1}{\epsilon^c})$. It would be interesting to see if there is any connection with the (still) open problem of finding the right bound on the size of ϵ-nets for the primal set system induced by lines in the plane.

References

[1] P. K. Agarwal, J. Pach, and M. Sharir. State of the union (of geometric objects): A review. In J. Goodman, J. Pach, and R. Pollack, editors, *Computational Geometry: Twenty Years Later*, pages 9–48. AMS, 2008.

[2] B. Aronov, E. Ezra, and M. Sharir. Small-size ϵ-nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing*, 39(7):3248–3282, 2010.

[3] S. Arya, G. D. da Fonseca, and D. Mount. Optimal area-sensitive bounds for polytope approximation. In *Proceedings of the International Symposium on Computational Geometry (SoCG)*, pages 363–372, 2012.

[4] I. Bárány. Random polytopes, convex bodies, and approximation. In *Stochastic Geometry*, pages 77–118. Springer, 2007.

[5] I. Bárány and D. G. Larman. Convex bodies, economic cap coverings, random polytopes. *Mathematika*, 35:274–291, 1988.

[6] H. Brönnimann, B. Chazelle, and J. Pach. How hard is half-space range searching. *Discrete & Computational Geometry*, 10:143–155, 1993.

[7] B. Chazelle. *The Discrepancy Method: Randomness and Complexity*. Cambridge University Press, 2000.

[8] C. Chekuri, K. L. Clarkson, and S. Har-Peled. On the set multi-cover problem in geometric settings. In *Proceedings of the International Symposium on Computational Geometry (SoCG)*, pages 341–350, 2009.

[9] G. Ewald, D. G. Larman, and C. A. Rogers. The directions of the line segments and of the r-dimensional balls on the boundary of a convex body in euclidean space. *Mathematika*, 17:1–20, 1970.

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

[11] J. Komlós, J. Pach, and G. Woeginger. Almost tight bounds for epsilon nets. *Discrete & Computational Geometry*, pages 163–173, 1992.

[12] A. Kupavskii, N. H. Mustafa, and J. Pach. New lower bounds for epsilon-nets. In *Proceedings of the International Symposium on Computational Geometry (SoCG)*, 2016.

[13] A. M. Macbeath. A theorem on non-homogeneous lattices. *Annals of Math*, 56:269–293, 1952.

[14] J. Matoušek. Efficient partition trees. *Discrete & Computational Geometry*, 8:315–334, 1992.

[15] J. Matoušek. Reporting points in halfspaces. *Computational Geometry*, 2(3):169 – 186, 1992.

[16] J. Matoušek. *Geometric Discrepancy: An Illustrated Guide*. Springer, 1999.

[17] J. Matoušek. *Lectures in Discrete Geometry*. Springer, 2002.

[18] N. H. Mustafa. A simple proof of the shallow packing lemma. *Discrete & Computational Geometry*, 55(3):739–743, 2016.

[19] N. H. Mustafa and S. Ray. Near-optimal generalisations of a theorem of Macbeath. In *Proceedings of the International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 578–589, 2014.

[20] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, 1995.

[21] J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. *Journal of the AMS*, 26:645–658, 2013.

[22] E. Pyrga and S. Ray. New existence proofs epsilon-nets. In *Proceedings of the International Symposium on Computational Geometry (SoCG)*, pages 199–207, 2008.