New Existence Proofs for ϵ-Nets

Evangelia Pyrga [*]
pyrga@mpi-inf.mpg.de

Saurabh Ray[†]
saurabh@cs.uni-sb.de

Abstract

We describe a new technique for proving the existence of small ϵ-nets for hypergraphs satisfying certain simple conditions. The technique is particularly useful for proving $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ upper bounds which is not possible using the standard VC dimension theory. We apply the technique to several geometric hypergraphs and obtain simple proofs for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for them. This includes the geometric hypergraph in which the vertex set is a set of points in the plane and the hyperedges are defined by a set of pseudo-disks. This result was not known previously. We also get a very short proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspaces in $\mathbb{R}^3$.

1 Introduction

A hypergraph is a pair (X, S) where X is a set called the *vertex set*, and S is a family of subsets of X called *hyperedges*. A finite hypergraph is a hypergraph whose vertex set is finite. Given a finite hypergraph (X, S), the general ϵ-net problem asks for a small set $Y \subseteq X$ such that for every hyperedge $s \in S$ containing more than $\epsilon|X|$ elements, $Y \cap s \neq \emptyset$. In a celebrated result, Haussler and Welzl [12] showed that if a hypergraph has VC dimension d, then picking a random sample of size $O(\frac{d}{\epsilon}\log\frac{d}{\epsilon})$ from X yields an ϵ-net with probability close to 1. Komlós et al. [13] proved that this is also tight. However, $O(\frac{1}{\epsilon})$ size ϵ-nets are known for various hypergraphs, typically hypergraphs induced by geometric objects. Pach and Woeginger [21] proved that halfspaces and translates of polytopes in $\mathbb{R}^2$ admit ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek et al. [18] proved that halfspaces in $\mathbb{R}^3$ and certain special families of pseudo-disks in $\mathbb{R}^2$ (they require that there is exactly one pseudo-disk through any three non-collinear points in the plane) admit ϵ-nets of size $O(\frac{1}{\epsilon})$. Matousek later found a shorter proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspaces in $\mathbb{R}^3$ via shallow cuttings [15]. Clarkson and Varadarajan [10] found a connection between the union complexity of a set of geometric objects and the size of ϵ-nets for the hypergraphs in which the objects form the vertex set and hyperedges are defined by common intersections of these objects. They show that if there is a canonical decomposition of the exterior of the union of any j of the objects into at most $f(j) \geq j$ simple regions then there is an ϵ-net of size $O(f(\frac{1}{\epsilon}))$. This gives a general technique for proving the existence of $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ size ϵ-nets for many hypergraphs induced by geometric objects. Still, there are many cases where their technique does not apply and we are still far from properly characterizing hypergraphs that admit ϵ-nets of size $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$. Apart from their inherent interest, ϵ-nets have several applications in computational geometry, statistics and learning theory (see

---

[*] Max-Planck Institut für Informatik

[†] Universität des Saarlandes

[16],[9],[19],[20],[4],[23]). Proving the existence of small ϵ-nets is also important since it implies better approximation algorithms for the corresponding set-cover problems (see [7], [10], [14]).

2 Our Contribution

As mentioned earlier, the VC dimension theory doesn't give $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ upper bounds on ϵ-net sizes and although $O(\frac{1}{\epsilon})$ size ϵ-nets are known for various special cases, no general technique for proving $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ size ϵ-nets is known. The technique used by Clarkson and Varadarajan [10] is very useful for hypergraphs in which the vertex set is a set of geometric objects and the hyperedges are defined by their common intersections. However, it is not clear how to extend the technique for general hypergraphs or even for geometric cases where the vertex set is a set of points and the hyperedges are defined by a set of geometric objects. We show that general hypergraphs satisfying certain simple conditions admit small ϵ-nets. The proof uses elementary techniques and suggests an algorithm to compute the ϵ-net (although we do not discuss algorithmic issues in this paper). We apply the technique to several hypergraphs induced by geometric objects and prove the existence $O(\frac{1}{\epsilon})$ size ϵ-nets for them. In particular, we get a very short proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspaces in $\mathbb{R}^3$. We also prove existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for hypergraphs in which the vertex set is a set of k-admissible regions in $\mathbb{R}^2$ (which includes pseudo-disks since they are 2-admissible) and the hyperedges are defined by their common intersections. We use this result to prove the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for the hypergraphs in which the vertex set is a set of points in $\mathbb{R}^2$ and the hyperedges are defined by a set of pseudo-disks. This result was not previously known and is an example of a case where the method used in [10] does not apply. Along the way, we also explore interesting properties of hypergraphs induced by geometric objects and points in the plane. For instance, we prove that given a set of points in the plane, the hypergraph induced by a set of k-admissible regions is 4-colorable.

3 Preliminaries

A hypergraph $\mathcal{H}$ is a pair $\mathcal{H} = (X, S)$ where X is a set called the *vertex set*, and S is a family of subsets of X called *hyperedges*. For any $Y \subseteq X$ we call $S|_Y = \{s \cap Y : s \in S\}$ the *projection* of S on Y. The projection of $\mathcal{H}$ on Y is the hypergraph $\mathcal{H}|_Y = (Y, S|_Y)$ (see [6] for a thorough treatment of hypergraphs). A set A is said to be *shattered* by S if all subsets of A can be obtained by intersecting A with members of S i.e. $S|_A = 2^A$. The *Vapnik-Chervonenkis dimension* (VC dimension in short) of $\mathcal{H}$, denoted by $dim(\mathcal{H})$, is the cardinality of the largest set $A \subseteq X$ shattered by S.

A finite hypergraph is a hypergraph with a finite vertex set. Let $\mathcal{H} = (X, S)$ be a finite hypergraph with $|X| = n$ and let $0 < \epsilon < 1$ be a parameter. An ϵ-net for $\mathcal{H}$ is a subset $Y \subseteq X$ s.t. $Y \cap s \neq \emptyset$ for every $s \in S$ with $|s| > \epsilon n$. We use the following result in this paper:

Theorem 1 (ϵ-net theorem [12]). *For any finite hypergraph $\mathcal{H} = (X, S)$ with $dim(\mathcal{H}) \leq d$ and parameters $0 < \epsilon < 1$, $\delta > 0$, a random subset $N \subseteq X$ of size $\max\left\{\frac{8d}{\epsilon}\log\frac{8d}{\epsilon}, \frac{4}{\epsilon}\log\frac{2}{\delta}\right\}$ is an ϵ-net for $\mathcal{H}$ with probability at least $1-\delta$.*

The hypergraphs we consider in this paper are induced by finite sets of points and geometric objects. Let P be a finite set of points and S a finite set of geometric objects. For an object $s \in S$, let $P(s)$ be the set of points contained in s, i.e., $P(s) = \{p \in P : p \in s\}$. Similarly for a point p, let $S(p)$ be the set of objects containing p, i.e., $S(p) = \{s \in S : p \in s\}$. The sets P and S induce two natural kinds of hypergraphs. If we treat the set of points P as the vertex set and the let the objects in S define the hyperedges, we get the hypergraph $(P, \{P(s) : s \in S\})$ which we call the *primal hypergraph induced by* P *and* S and denote it by $\mathcal{H}(P, S)$. On the other hand, if we think of the set of objects S as the vertex set and let the points in P define the hyperedges, we get the *dual hypergraph induced by* P *and* S, denoted by $\mathcal{H}^*(P, S) = (S, \{S(p) : p \in P\})$.

4 Halfspaces in $\mathbb{R}^3$

In this section we give a simple proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for halfspaces in $\mathbb{R}^3$. This result was first proved in [18] and later a simpler proof appeared in [15]. For convenience, we consider the ϵ-net problem for the dual hypergraph induced by finite sets of points and halfspaces in $\mathbb{R}^3$. The existence of ϵ-nets of size $O(\frac{1}{\epsilon})$ for such dual hypergraphs implies the same for the primal hypergraphs since the roles of points and halfspaces can be exchanged by using projective duality (see [11]).

Given that the dual hypergraph induced by a set of points and a set of halfspaces in $\mathbb{R}^3$ has a finite VC dimension, it follows from the ϵ-net theorem (Theorem 1) that such a hypergraph admits an ϵ-net whose size depends only on the parameter ϵ. In other words, when ϵ is a constant, there is an ϵ-net of constant size. We use this fact and the following claim to prove that the dual hypergraph induced by a given finite set P of points and a set H of halfspaces in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.

Claim 1. *Given any finite set* Q *of points in* $\mathbb{R}^3$*, there exists a graph* $G_Q = (Q, E_Q)$ *with at most* $4|Q|$ *edges such that for any halfspace* h *in* $\mathbb{R}^3$*, the subgraph of* G_Q *induced by the points of* Q *contained in* h *(i.e.* $Q \cap h$*) is connected.*

Proof. We construct G_Q as follows: Let $Q' \subseteq Q$ be the vertices of the convex hull $\mathrm{CH}(Q)$ of Q. We include the edges of the 1-skeleton of $\mathrm{CH}(Q)$ (i.e. the graph with the vertices of $\mathrm{CH}(Q)$ as the vertex set and the edges (1-faces) as the edge set) in G_Q. For each of the points $q \in Q \backslash Q'$, we pick a tetrahedron containing q whose vertices are in Q' (there is always such a tetrahedron by Carathéodory's theorem [17]) and put edges between q and each of the four corners of this tetrahedron. The construction of G_Q is complete. It contains at most $4|Q|$ edges since the 1-skeleton of $\mathrm{CH}(Q)$ is a planar graph and each point in $Q \backslash Q'$ has degree four. For any halfspace $h \in \mathbb{R}^3$, the subgraph of G_Q induced by the points in $Q' \cap h$ is obviously connected and each point in $(Q \backslash Q') \cap h$ is connected to at least one of the points of $Q' \cap h$. Therefore, the subgraph induced by the points in $Q \cap h$ is connected. $\square$

Let $n = |H|$. We say that a point $p \in P$ is *dense* if it is contained in more than ϵn of the halfspaces in H. We further classify the dense points into *highly dense* and *moderately dense* depending on whether they are contained in more than or at most $2\epsilon n$ of the halfspaces in H, respectively. An ϵ-net for the dual hypergraph $\mathcal{H}^*(P, H)$ is a subset $Y \subseteq H$ which *covers* all dense points i.e. each of the dense points is contained in at least one of the halfspaces in Y. We say that a subset $Z \subseteq H$ is a *moderate* ϵ-net for $\mathcal{H}^*(P, H)$ if it covers all the moderately dense points in P.

Claim 2. *The hypergraph* $\mathcal{H}^*(P, H)$ *admits a moderate* ϵ*-net of size* $O(\frac{1}{\epsilon})$.

Proof. Let $M \subseteq P$ be the set of moderately dense points in P. For each $p \in M$, let $H(p)$ denote the set of halfspaces in H which contain p. We say that two points $p, q \in M$ are *independent* if $|H(p) \cap H(q)| \leq \epsilon n/8$. Let $I \subseteq M$ be an inclusion-maximal set of pairwise independent points in M. The maximality of I implies that for any $p \in M$, there is a point $q \in I$, not necessarily different from p, such that $|H(p) \cap H(q)| > \epsilon n/8$. Since q is a moderately dense point, $H(q) \leq 2\epsilon n$ and hence $|H(p) \cap H(q)| > |H(q)|/16$. This means that a $\frac{1}{16}$-net for $\mathcal{H}^*(P, H(q))$ covers p. In other words, $Z = \bigcup_{r \in I} Y_r$, where Y_r denotes a $\frac{1}{16}$-net for $\mathcal{H}^*(P, H(r))$, is a moderate ϵ-net for $\mathcal{H}^*(P, H)$. The size of such a net is $O(t)$, where $t = |I|$, since each of the $\frac{1}{16}$-nets is of constant size. We now show that $t = O(\frac{1}{\epsilon})$.

By Claim 1, there is a graph $G_I = (I, E_I)$ such that $|E_I| \leq 4t$ and for any $h \in H$, $|h \cap I|$ induces a connected subgraph of G_I. We say that a halfspace h contains an edge $e \in E_I$ if both the endpoints of e are contained in h. We denote the set of halfspaces containing an edge e by $H(e)$. For any halfspace h, let $n_h = |h \cap I|$ and let m_h be the number of edges contained in h. Since $h \cap I$ induces a connected subgraph of G_I, $n_h - m_h \leq 1$ for each h. Summing over the n halfspaces in H,

$$\sum_{h \in H} n_h - \sum_{h \in H} m_h \leq n. \tag{1}$$

Now, since each $p \in I$ is a dense point,

$$\sum_{h \in H} n_h = \sum_{p \in I} |H(p)| > t\epsilon n. \tag{2}$$

Each edge in E_I is contained in at most $\epsilon n/8$ halfspaces since both its endpoints belong to I. Therefore,

$$\sum_{h \in H} m_h = \sum_{e \in E_I} |H(e)| \leq |E_I| \frac{\epsilon n}{8} \leq 4t \frac{\epsilon n}{8}. \tag{3}$$

It follows from (1),(2) and (3) that $t \leq 2/\epsilon$ and hence Z is a moderate ϵ-net of size $O(\frac{1}{\epsilon})$. ☐

Theorem 2. *The dual hypergraph $\mathcal{H}^*(P, H)$ induced by a set of points P and a set of halfspaces H in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $M \subseteq H$ be a 2ϵ-net for $\mathcal{H}^*(P, H)$ and Z a moderate ϵ-net for $\mathcal{H}^*(P, H)$. Then $Z \cup M$ is an ϵ-net for $\mathcal{H}^*(P, H)$ since Z covers all of the moderately dense points and M covers all the highly dense points. By Claim 2, there exists a moderate ϵ-net of size $O(1/\epsilon)$. If we denote the size of the smallest ϵ-net admitted by $\mathcal{H}^*(P, H)$ by $f(\epsilon)$, we have

$$\begin{aligned} &f(x) = 0, \quad \forall\ x \geq 1, \\ &f(\epsilon) \leq O(1/\epsilon) + f(2\epsilon). \end{aligned}$$

It follows that $f(\epsilon) = O(\frac{1}{\epsilon})$. ☐

Using projective duality between points and halfspaces, we also obtain the next theorem.

Theorem 3. *The primal hypergraph $\mathcal{H}(P, H)$ induced by a finite set of points P and a finite set of halfspaces H in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

5 Abstract Framework

The proof for the existence of $O(\frac{1}{\epsilon})$ size ϵ-nets for hypergraphs induced by points and halfspaces in R^3 can adapted for any hypergraph $\mathcal{H} = (X, S)$ which has the properties that we have exploited in the proof. In the following, we denote the set of hyperedges containing a particular element $x \in X$ by $S(x)$.

Theorem 4. *Any hypergraph $\mathcal{H} = (X, S)$ satisfying the following two conditions admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

1. *For any $0 < \epsilon < 1$ and for any $Y \subseteq X$, $\mathcal{H}|_Y$ admits an ϵ-net whose size depends only on ϵ.*
2. *There exist constants $\alpha > 0, \beta \geq 0$ and $\tau > 0$ s.t. for any $I \subseteq S$, there is a graph $G_I = (I, E_I)$ with $|E_I| \leq \beta |I|$ so that for any element $x \in X$ we have $m_x \geq \alpha n_x - \tau$, where $n_x = |I(x)|$ and m_x is the number of edges in E_I whose both endpoints (which are hyperedges) contain x.*

Proof. The proof is analogous to the proof for halfspaces in $\mathbb{R}^3$. As before, we call a hyperedge $s \in S$ *dense* if $|s| > \epsilon n$, where $n = |X|$. We say that a dense hyperedge s is *highly dense* if $|s| > 2\epsilon n$ and *moderately dense* otherwise. An ϵ-net for $\mathcal{H}$ is a subset $Y \subseteq X$ which *covers* all dense hyperedges i.e. each dense hyperedge contains at least one element of Y. A moderate ϵ-net is a subset $Z \subseteq X$ which covers all moderately dense hyperedges.

We show that $\mathcal{H}$ admits a moderate ϵ-net of size $O(\frac{1}{\epsilon})$. From this we can conclude the existence of an $O(\frac{1}{\epsilon})$ size ϵ-net for $\mathcal{H}$ by an argument analogous to the proof of Theorem 2. We say that two hyperedges s and s' are *independent* if $|s \cap s'| \leq \frac{\alpha}{2\beta}\epsilon n$. Let $I \subseteq S$ be a inclusion-maximal set of pairwise independent moderately dense hyperedges. Then for each $s' \in S$, there is a set $s \in I$ such that $|s \cap s'| > \frac{\alpha}{2\beta}\epsilon n$ which implies that $|s \cap s'| > \frac{\alpha}{4\beta}|s|$ since $|s| \leq 2\epsilon n$. Therefore, an $\frac{\alpha}{4\beta}$-net for $\mathcal{H}_s$ covers s'. This means that $Z = \bigcup_{r \in I} Y_r$, where Y_r denotes a $\frac{\alpha}{4\beta}$-net for $\mathcal{H}|_r$, is a moderate ϵ-net for $\mathcal{H}$. Moreover, $|Z| = O(t)$ where $t = |I|$, since each of the $\frac{\alpha}{4\beta}$-nets are of constant size due to the first condition in the statement of the theorem. Now we show that $t = O(\frac{1}{\epsilon})$.

Let $G_I = (I, E_I)$ be the graph ensured by the second condition in the statement of the theorem. For an edge $e = (s, s') \in E_I$, let $X(e) = |s \cap s'|$. Since the hyperedges in I are pairwise independent, $X(e) \leq \frac{\alpha}{2\beta}\epsilon n$. For each $x \in X$, we have $\alpha n_x - m_x \leq \tau$, where $n_x = |I(x)|$ and m_x is the number of egdes in E_I whose both endpoints contain x. Summing over all $x \in X$, we have:

$$\sum_{x \in X} \alpha n_x - \sum_{x \in X} m_x \leq \tau n. \tag{4}$$

Now,

$$\sum_{x \in X} n_x = \sum_{s \in I} |s| \geq t\epsilon n \tag{5}$$

since each $s \in I$ is dense. Also, since $X(e) \leq \frac{\alpha}{2\beta}\epsilon n$ for each $e \in E_I$ and $|E_I| \leq \beta t$,

$$\sum_{x \in X} m_x = \sum_{e \in E_I} X(e) \leq |E_I| \frac{\alpha}{2\beta}\epsilon n \leq \beta t \frac{\alpha}{2\beta}\epsilon n = \frac{\alpha}{2} t\epsilon n. \tag{6}$$

From (4), (5) and (6) we get:

$$\alpha t\epsilon n - \frac{\alpha}{2} t\epsilon n \leq \tau n \implies t \leq \frac{2\tau}{\alpha} \cdot \frac{1}{\epsilon}. \tag{7}$$

Since α and τ are constants, $t = O(\frac{1}{\epsilon})$. Hence, Z is a moderate ϵ-net of size $O(\frac{1}{\epsilon})$ for $\mathcal{H}$ and we conclude from a calculation similar to the one in the proof of Theorem 2 that $\mathcal{H}$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$. $\square$

The second condition of Theorem 4 requires $|E_I|$ to be $O(|I|)$. It is natural to expect that if instead we had $|E_I| \leq |I| b(|I|)$ for some *small* function $b(\cdot)$ then we should still be able to prove the existence of a correspondingly *small* ϵ-net. Indeed, this is true and the following theorem can be proved along the lines of the proof of Theorem 4.

Theorem 5. *Let $\mathcal{H} = (X, S)$ be a hypergraph satisfying the following two conditions:*

1. *For any ϵ and for any $Y \subseteq X$, $\mathcal{H}|_Y$ admits an ϵ-net whose size depends only on ϵ.*
2. *There exist constants $\alpha > 0, \tau > 0$ and a positive non-decreasing sublinear function $b(\cdot)$ s.t. for any $I \subseteq S$, there is a graph $G_I = (I, E_I)$ with $|E_I| \leq |I|\, b(|I|)$ so that for any element $x \in X$ we have $m_x \geq \alpha n_x - \tau$, where $n_x = |I(x)|$ and m_x is the number of edges in E_I whose both endpoints (which are hyperedges) contain x.*

Then, $\mathcal{H}$ admits an ϵ-net of size $O\left(\frac{1}{\epsilon} \cdot (4\tau)^{f^*\left(\frac{2}{\epsilon}\right)} \tilde{f}\left(\frac{2}{\epsilon}\right)\right)$ where $f^*(\cdot)$ and $\tilde{f}(\cdot)$ are defined as:

$$f^*(k) = \begin{cases} 0, & \textit{if } f(k) \geq k \\ 1 + f^*(f(k)), & \textit{otherwise} \end{cases}$$

$$\tilde{f}(k) = \begin{cases} 1, & \textit{if } f(k) \geq k \\ f(k) \cdot \tilde{f}(f(k)), & \textit{otherwise} \end{cases}$$

The proof Theorem 4 can be easily adapted for the case in which the vertices have positive weights (instead of all vertices having weight 1). The same holds for Theorem 5. The proof of Theorem 4 also suggests an algorithm for computing an ϵ-net of size $O(\frac{1}{\epsilon})$. It can, in fact, be turned into an efficient algorithm when the VC dimension of the hypergraph is finite. However, we do not discuss these issues in this paper.

6 Geometric Applications

In the following, we present several applications of Theorem 4. The geometric hypergraphs that we consider here have finite VC dimension and hence automatically satisfy the first condition of Theorem 4 (due to Theorem 1). Hence, we only prove that they satisfy the second condition and conclude the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for them. Most of the results presented here have been proven before using various techniques. Apart from Theorems 10 and 12, which were not previously known, the remaining results also follow from the framework of Clarkson and Varadarajan [10]. We include them here in order to demonstrate that they follow from our framework too. Also, in some cases, our technique leads to a simpler proof. Theorems 10 and 12 show a way to overcome the limitations of the technique used in [10].

Most of the definitions given in this section are more thoroughly explained in [1] and [2].

6.1 Translates of Orthants in $\mathbb{R}^3$

We will show that the dual hypergraph induced by a finite set of points and translates of an orthant in $\mathbb{R}^3$ (also called an octant) admits an ϵ-net of size $O(\frac{1}{\epsilon})$. Let P be a finite set of points in $\mathbb{R}^3$ and let O be the set of all translates of some orthant in $\mathbb{R}^3$. We will also denote a point $p \in \mathbb{R}^3$ as (x_p, y_p, z_p), where x_p, y_p and z_p are the x, y and z coordinates of p. For $p, q \in \mathbb{R}^3$, we will write $p \geq q$ iff $x_p \geq x_q, y_p \geq y_q$ and $z_p \geq z_q$. We define the notation $p \leq q$ in a similar manner. W.l.o.g. we assume that the orthants in O are axis-parallel, and every $T \in O$ is of the form $\{(x, y, z) : x \geq x_T, y \geq y_T, z \geq z_T\}$, for some $(x_T, y_T, z_T) \in \mathbb{R}^3$, which we call the *corner* of T. For any $T \in O$ and any $Q \subseteq P$, let $Q(T) = Q \cap T$. Note that if an orthant $T \in O$ contains some $p \in Q$, then it also contains any point $q \in Q$ such that $q \geq p$. For any two points $p, q \in Q$, we define the *minimal common orthant of* p *and* q, $T_{p,q} \in O$, as the minimal (w.r.t. inclusion) orthant that contains both p and q. The corner of $T_{p,q}$ is the point $(\min\{x_p, x_q\}, \min\{y_p, y_q\}, \min\{z_p, z_q\})$.

Lemma 1. *For any $Q \subseteq P$, there is a graph $G_Q = (Q, E_Q)$, such that $|E_Q| \leq 3|Q|$ and for any $T \in O$, there are at least $\frac{1}{2}|Q(T)| - 1$ edges among the points in $Q(T)$.*

Proof. For a point $p \in Q$ we define the *x-neighbor* of p as $N_x(p) = \arg\max_q\{x_q : y_q \geq y_p, z_q \geq z_p, q \in Q \setminus \{p\}\}$. The y- and z-neighbors $N_y(p)$ and $N_z(p)$ are defined similarly. Note that it is not necessary that all three of $N_x(p), N_y(p), N_z(p)$ exist for all $p \in Q$. For every $p \in Q$, we add to E_Q the edges $(p, N_x(p)), (p, N_y(p))$ and $(p, N_z(p))$, whenever $N_x(p), N_y(p)$ and $N_z(p)$ exist, respectively. The construction of G_Q is now complete. Clearly, $|E_Q| \leq 3|Q|$, since every point in Q accounts for at most 3 edges.

Consider any orthant $T \in O$. We claim that there is at most one point $p \in Q(T)$ that does not share an edge with another point in $Q(T)$. This will immediately imply that the number of edges whose both endpoints are contained in Q is at least $\frac{1}{2}|Q(T)| - 1$. Assume, for contradiction, that there are $p, q \in Q(T)$, $p \neq q$, neither of which shares an edge with another point in $Q(T)$. Since $p, q \in T$, their minimal common orthant $T_{p,q}$ is also contained in T. W.l.o.g., assume that the corner of $T_{p,q}$ is the point $o_{p,q} = (x_p, y_p, \min\{z_p, z_q\})$, i.e. at least two coordinates of the corner (namely the x and y coordinates) are equal to the corresponding coordinates of p (any other case can be treated similarly). Consider the set of points $Z_p = \{q' \in Q \setminus \{p\} : x_{q'} \geq x_p, y_{q'} \geq y_p\}$. The z-neighbor of p is given by $N_z(p) = \arg\max_{q'}\{z_{q'} : q' \in Z_p\}$. Since $q \in Z_p(T)$, it must be that $z_q \leq z_{N_z(p)}$, implying that $N_z(p) \geq o_{p,q}$, i.e. $N_z(p) \in T$. But then the edge $(p, N_z(p)) \in E_Q$ contradicting our assumption that p does not share an edge with another point in $Q(T)$. Therefore, there is at most one point in $Q(T)$ which does not share an edge with another point in $Q(T)$, thus proving the claim. $\square$

Lemma 1 and Theorem 4 imply the following theorem:

Theorem 6. *The dual hypergraph defined by a finite set of points and a finite set of translates of an orthant in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

The above theorem also implies the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for the primal hypergraph. To see this, note that we can substitute every orthant T by its corner o_T, and every point $p = (x_p, y_p, z_p)$ by an orthant of the form $\{(x, y, z) : x \leq x_p, y \leq y_p, z \leq z_p\}$. This preserves the incidences between orthants and points. Therefore, we have:

Theorem 7. *The primal hypergraph defined by a finite set of points and a finite set of translates of an orthants in $\mathbb{R}^3$ admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

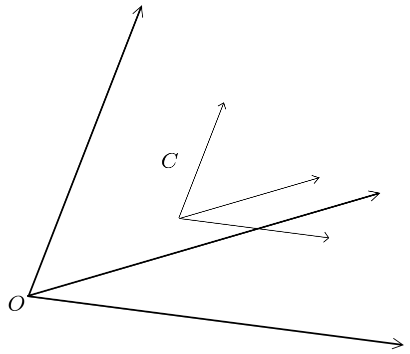

Figure 1: Cone C behaves like an orthant in the oblique coordinate system with origin O.

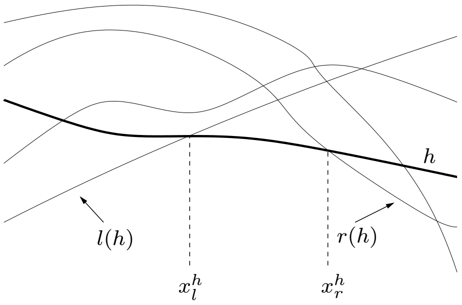

Figure 2: $l(h)$ and $r(h)$ for the pseudo-halfplane h in a family of pseudo-halfplanes.

Remark 1. *Since cones with a triangular cross-section behave like orthants in a suitable oblique coordinate system (see Fig. 1), the same result follows for translates of such* triangular *cones in* $\mathbb{R}^3$*. From this, it is easy to show that translates of a tetrahedron* Δ *admit an* ϵ*-net of size* $O(\frac{1}{\epsilon})$*. We only need to consider a fine enough (oblique) gridding of* $\mathbb{R}^3$ *so that every translate of* Δ *has at most one corner in any cell and at the same time it does not intersect more than a constant number of cells. Then, inside any cell, we can treat the translates of* Δ *intersecting it as cones. Since a convex polytope in* $\mathbb{R}^3$ *with* k *vertices can be triangulated with* $O(k)$ *tetrahedra, it follows that translates of a convex polytope in* $\mathbb{R}^3$ *with* k *vertices admit an* ϵ*-net of size* $O(\frac{k^2}{\epsilon})$*.*

6.2 Pseudo-Halfplanes in $\mathbb{R}^2$

A family of (x-monotone) *pseudo-lines* in the plane is a set of graphs of continuous univariate functions, which intersect in at most one point and cross at that point. A family of *pseudo-halfplanes* is a set of closed sets in the plane whose boundaries form a family of pseudo-lines. For convenience, we will just write *halfplanes* for pseudo-halfplanes. For any halfplane h, we will denote the function tracing its boundary by f_h. With a slight abuse of notation, we will also refer to the boundary of h by f_h. A family of *upper* halfplanes is a set of halfplanes each of which is bounded from below, i.e. for each halfplane h in the set, $h = \{(x, y) \in \mathbb{R}^2 : y \geq f_h(x)\}$. Similarly, a family of *lower* halfplanes is a set of halfplanes each of which is bounded from above.

Let $\mathscr{H}$ be a finite family of upper halfplanes and $X \subseteq \mathscr{H}$. For any $p \in \mathbb{R}^2$, we will denote the set of halfplanes in X that contain p by $X(p)$, i.e., $X(p) = \{h \in X : p \in h\}$.

Lemma 2. *For any $X \subseteq \mathscr{H}$, there is a graph $G_X = (X, E_X)$, such that $|E_X| \leq 6|X|$ and for any $p \in \mathbb{R}^2$, the halfplanes in $X(p)$ induce a connected subgraph (which therefore contains at least $|X(p)| - 1$ edges).*

Proof. For halfplanes h and h', we say that h lies *below* h' at x, if $f_{h'}(x) < f_h(x)$. Note that if a point $p = (x, y)$ is contained in some $h \in X$, then p is also contained in every h' lying below h at x. Therefore, in order to construct G_X in such a way that the set of halfplanes containing p induce a connected subgraph, it suffices to ensure that for all $x \in \mathbb{R}$, each halfplane $h \in X$ shares an edge with some halfplane $h' \in X$ below it at x (if one such h' exists). For simplicity, we assume that the boundaries f_h and $f_{h'}$, of any pair of halfplanes $h, h' \in X$, cross exactly once. (If there are h, h' such that the boundaries f_h and $f_{h'}$ never cross, then for all $x \in \mathbb{R}$, one of them, say h, lies above the other. Therefore, we can put an edge between h and h' and ignore pseudo-halfplane h.) For a halfplane $h \in X$, let $l(h)$ be the halfplane in X which lies below h for the maximal interval (w.r.t. inclusion) of the form $(-\infty, x)$. Similarly let $r(h)$ be the halfplane of X which lies below h for the maximal interval of the form $(x, +\infty)$. More formally,

$$\begin{aligned} l(h) \quad &= \arg\max\nolimits_{h' \in X} \{\, x_l : \quad f_{h'}(x_l) = f_h(x_l) \text{ and} \\ &\qquad\qquad\qquad\qquad\quad \forall x < x_l,\ f_{h'}(x) < f_h(x)\ \}\,, \\ r(h) \quad &= \arg\min\nolimits_{h' \in X} \{\, x_r : \quad f_{h'}(x_r) = f_h(x_r) \text{ and} \\ &\qquad\qquad\qquad\qquad\quad \forall x > x_r,\ f_{h'}(x) < f_h(x) \}\,. \end{aligned}$$

Let x^l_h and x^r_h be the x-coordinates of the intersection of f_h with $f_{l(h)}$ and $f_{r(h)}$, respectively. For an example, see Fig. 2.

We construct the edge set E_X in two stages. In the first stage, we connect every $h \in X$ with $l(h)$ and $r(h)$ (unless these edges already exist). This way, we add at most $2|X|$ edges and every $h \in X$ is connected to another halfplane that lies below it, for all $x \in (-\infty, x^h_l) \cup (x^h_r, +\infty)$. In the second stage, for every $h \in X$, we restrict our attention only to the interval $I_h = [x^l_h, x^r_h]$. (If $x^l_h > x^r_h$, we can completely ignore h for this stage, since the edges added to it in the first stage suffice. If x^l_h or x^r_h does not exist, we can consider it equal to $-\infty$ or $+\infty$, respectively.) Let σ_h denote the drawing of f_h restricted to I_h. We claim that the x-monotone curves σ_h (for all $h \in X$) form a set Γ of curves whose interiors do not cross. For contradiction, assume that the interiors of σ_{h_1} and σ_{h_2} cross at a point (x_0, y_0), and (w.l.o.g. assume that) for $x < x_0$, σ_{h_1} lies below σ_{h_2} (i.e., $f_{h_1}(x) < f_{h_2}(x)$). Then, either $x^{h_2}_l \geq x_0$ or $x^{h_1}_r \leq x_0$, contradicting the assumption that the interiors of σ_{h_1} and σ_{h_2} intersect at (x_0, y_0). For each $h \in X$ and $x \in I_h$, we will ensure that h is connected to some $h' \in X$ such that $x \in I_{h'}$ and $\sigma_{h'}$ lies below σ_h at x (if such an h' exists). To achieve this, we consider a trapezoidal decomposition of the segments σ_h (as in [11], Chapter 6, noting that having curved segments instead of line segments doesn't cause any problems). The decomposition consists of at most $3|X| + 1$ non-overlapping trapezoids, whose upper and lower sides are parts of curves in Γ and the left and right sides are vertical line segments. For each trapezoid, we put an edge between the halfplanes corresponding to the upper and lower sides. The construction of G_X is now complete.

To see that for any $p \in \mathbb{R}^2$ the subgraph of G_X induced by the halfplanes in $X(p)$ is connected, observe first that the lower envelope of the arrangement of halfplanes in X is identical to the lower envelope of the segments in Γ. Now, for each $h \in X$ and $x \in \mathbb{R}$: If $x \in (-\infty, x^h_l) \cup (x^h_r, +\infty)$, then h is connected to some halfplane lying below it at x (if one such exists) due to the edges added in the first stage. Otherwise $x \in I_h$, and the only case in which h is not connected to any halfplane lying below it is when σ_h is the lowest segment of Γ at x. However, in that case, h is also the lowest halfplane of X at x. Therefore for every $x \in \mathbb{R}$, each halfplane $h \in X$ shares an edge with some $h' \in X$, such that $f_{h'}(x) < f_h(x)$, unless h is the lowest halfplane at x, implying thus that the subgraph is connected. Since the total number of edges added in the two stages is at most $5|X| + 1 \leq 6|X|$, the Lemma follows. ☐

It follows from Theorem 4 and Lemma 2 that the primal hypergraph defined by a family of upper (or lower) pseudo-halfplanes and a set of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.

Theorem 8. *The primal hypergraph induced by a finite family of pseudo-halfplanes $\mathscr{H}$ and a finite set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $\mathscr{H}_l$ and $\mathscr{H}_u$ be the sets of lower and upper halfplanes (respectively) in $\mathscr{H}$. We construct separate ϵ-nets N_l and N_u for $\mathcal{H}(P, \mathscr{H}_l)$ and $\mathcal{H}(P, \mathscr{H}_u)$. Then, $N_l \cup N_u$ gives an ϵ-net of size $O(\frac{1}{\epsilon})$ for $\mathcal{H}(P, \mathscr{H})$. ☐

Using the duality between a family of pseudo-lines and a set of points in the plane, as defined in [3], we can exchange the roles of points and upper halfplanes in Lemma 2 and prove the following:

Theorem 9. *The dual hypergraph induced by a finite family of pseudo-halfplanes $\mathscr{H}$ and a finite set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

6.3 Pseudo-Parabolic Halfplanes in $\mathbb{R}^2$

A family of *pseudo-parabolas* is a set of graphs of continuous univariate functions every two of which intersect (and cross) in at most two points. (We assume that every tangency is equivalent to two intersections.) A family of *pseudo-parabolic halfplanes* (parabolic halfplanes for short) is a set of closed regions in the plane whose boundaries form a family of pseudo-parabolas. For a parabolic halfplane b, let f_b denote the function that defines the corresponding pseudo-parabola. We define *upper* and *lower* parabolic halfplanes just as we did for pseudo-halfplanes in Section 6.2.

Lemma 3. *Let $\mathscr{B}$ be a family of upper pseudo-parabolic halfplanes and P a set of points in $\mathbb{R}^2$. Then, for any $B \subseteq \mathscr{B}$ there is a graph $G_B = (B, E_B)$, such that $|E_B| \leq 6|B|$ edges, and for any $p \in P$ the parabolic halfplanes containing p induce a connected subgraph.*

Proof. We will assume that any $b_1, b_2 \in B$ cross exactly twice. It is not too hard to prove that for any $b_1, b_2 \in B$ that intersect only once, additional crossings can be created at the right of the rightmost point in P, without changing the incidences between the parabolic halfplanes and the points in P. (Again, if $b, b' \in B$ intersect tangentially or never intersect, then one of them, let it be b, lies completely above the other. Therefore, we can connect b to b' with an edge and ignore b from then on.) For each $b \in B$, we define $l(b), r(b), x_l^b, x_r^b$ and σ_b just as in the proof of Lemma 2. Note that the definitions imply that x_l^b is the first point of intersection between b and $l(b)$, while x_r^b is the second point of intersection between $r(b)$ and b. The segments σ_b again form a set of x-monotone curves whose interiors do not cross. Assuming the contrary, say that the interiors of σ_{b_1} and σ_{b_2} intersect at a point (x_0, y_0). If this is the first intersection between the curves f_{b_1} and f_{b_2}, and for $x < x_0$, $f_{b_1}(x) < f_{b_2}(x)$, then $x_l^{b_2} \geq x_0$, contradicting the assumption that the interiors of σ_{b_1} and σ_{b_2} intersect at (x_0, y_0). Similarly, (x_0, y_0) cannot be the second point of intersection between the two curves either. The rest of the proof is identical to the proof of Lemma 2. ☐

The next theorem follows from Lemma 3 and Theorem 4.

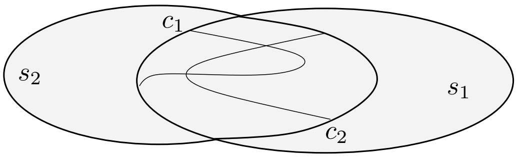

Figure 3: A chord of s_1 lies inside s_2 and vice-versa.

Theorem 10. *The primal hypergraph induced by a finite family of pseudo-parabolic halfplanes* $\mathscr{B}$ *and a set of points* P *in the plane admits an* ϵ*-net of size* $O(\frac{1}{\epsilon})$.

Unfortunately, there is no duality (similar to the one used for pseudo-lines) known for a set of pseudo-parabolas and a set of points in $\mathbb{R}^2$. Therefore, a different technique is required in order to prove the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ for the dual hypergraph defined by a family of pseudo-parabolic halfplanes and a set of points in the plane. We will, in fact, derive it from a more general result proved in the next section (see Remark 3).

6.4 k-admissible Regions in $\mathbb{R}^2$

Consider now a family of regions in $\mathbb{R}^2$, each of which is bounded by a closed Jordan curve. We will call it a family of *k-admissible* regions (for k even), if for any two s_1, s_2 of the regions, the Jordan curves bounding them cross (two Jordan curves cross when at a certain point a curve passes from one side of another curve to the other) in $l \leq k$ points, (for some even l), and both $s_1 \setminus s_2$ and $s_2 \setminus s_1$ are connected regions. A family of 2-admissible regions is also called a family of *pseudo-disks*.

Let S be a family of k-admissible regions, and P a *finite* set of points in $\mathbb{R}^2$. For any $Q \subseteq P$, we will show that there is a plane multigraph G (a crossing-free drawing of planar graph which may contain multiple edges between two vertices), such that the subgraph of G induced by the set of points contained in any of the regions $s \in S$, is connected. The graph will be given as the union of *connecting graphs* for each $s \in S$: For a region $s \in S$, we call a plane connected graph $G_s = (Q(s), E_s)$, where $Q(s) = Q \cap s$, an *s-connecting graph* if the drawing of the edges in E_s is strictly contained in s. Moreover, we say that a set of edges E *properly connects* a region $s \in S$, if there is a subset $E' \subseteq E$, such that the graph $G' = (Q(s), E')$ is an s-connecting graph.

In the following, whenever we refer to an edge, we will also refer to its drawing. We say that an edge e *pierces* a region s, if $s \setminus e$ has at least two disconnected components and not all the points of $Q \cap s$ lie in the same component. A *chord* of a region s is a Jordan arc with the endpoints lying on the boundary ∂s of s and the interior lying in the strict interior of s. If c is a chord of s, then $s \setminus c$ consists of exactly two disconnected regions.

Lemma 4. *Let* $s_1, s_2 \in S$. *Let* c_1 *be a chord of* s_1 *that lies in the interior of* s_2 *and* c_2 *be a chord of* s_2 *that lies in the interior of* s_1. *Then* c_1 *and* c_2 *cross at an even number of points.*

Proof. Fig. 3 shows a simple example which gives the intuition behind the lemma. We now prove it formally.
Since c_1 is a chord of s_1, it splits s_1 into two parts A and B. The chord c_1 lies in the interior of s_2 and $s_1 \setminus s_2$ has at most one connected component. Therefore, exactly one of A, B is contained in the interior of s_2. (If both A, B lie in s_2 then s_1 does not contain any point of ∂s_2, and hence it cannot contain a chord of s_2 either.) Assume that B is not contained in s_2. Since c_2 is a chord

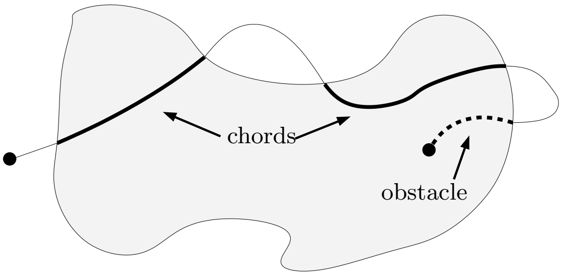

Figure 4: Chord and obstacle segments of an edge w.r.t. a region.

of s_2, the endpoints of c_2 are on ∂s_2. Also, A is contained in the strict interior of s_2 and therefore does not contain any points of ∂s_2. Hence, both endpoints of c_2 must lie in B, i.e. they are both on the same side of c_1 inside s_1. This immediately implies that c_1 and c_2 may only cross an even number of times. □

The following lemma indicates how to construct the aforementioned plane multigraph.

Lemma 5. *For any* $S' \subseteq S$ *and any given set of pairwise non-crossing* "compulsory" *edges* E_c*, such that no edge* $e \in E_c$ *pierces any of the regions in* S'*, there is a plane multigraph graph* $G = (Q, \mathbb{E} \cup D)$*, such that* $E_c \cup D$ *properly connects every* $s \in S'$*.*

Proof. Let $S' = \{s_1, s_2, \cdots, s_d\}$. We will use induction on the cardinality d of S'.

For $d = 1$, let I_1 be the set of points in the plane contained in the interior of s_1 and which do not lie in the interior of any of the edges in E_c. Since no edge in E_c pierces s_1, all points in $Q(s)$ belong to a connected component I'_1 of I_1. Therefore, there is a plane multigraph $G_1 = (Q, E_c \cup D_1)$, such that the edges in D_1 are strictly contained in I'_1 (and therefore in the strict interior of s_1) and $E_c \cup D_1$ properly connects s_1.

Assume now that for $d = l \geq 1$ and any compulsory set of pairwise non-crossing edges E_c which do not pierce any of $s \in S'$ there is a plane multigraph $G_l = (Q, E_c \cup E_l)$, with $E_c \cup E_l$ properly connecting $s_1, s_2, .., s_l$. For $d = l + 1$, let E'_l be the subset of edges in E_l that do not pierce s_{l+1}. Any edge $e \in E_l \setminus E'_l$ is split by ∂s_{l+1} into a set of segments. The segments that are contained in s and are not chords of s_{l+1} will be called *obstacles* (see for example Fig. 4). Note that one endpoint of each obstacle lies on ∂s_{l+1} and the other is one of the endpoints of the edge containing that obstacle. (Assuming general position, no point in Q, and therefore no edge endpoint, lies on the boundary of any region.) Let I_{l+1} be the set of points in $\mathbb{R}^2$ which are contained in the interior of s_{l+1} and which do not lie in the interior of any of the edges in $E_c \cup E'_l$ or the interior of any of the obstacles. Note that no edge in $E_c \cup E'_l$ or an obstacle pierces s_{l+1}. Moreover, since no two obstacles cross, any common point two of them may share will be an edge endpoint belonging to $I_{l+1} \cap Q(s_{l+1})$. Thus, all points in $Q(s_{l+1})$ belong to a connected component I'_{l+1} of I_{l+1}. Therefore, there is a set D' of edges contained in I'_{l+1} such that the plane multigraph $G_{l+1} = (Q, E_c \cup E'_l \cup D')$ properly connects s_{l+1}.

We claim that no edge $e \in D'$ pierces any of $s_1, \cdots, s_l$. For contradiction, assume that some $e \in D'$ pierces some s_i, $i \leq l$. Then, e contains a chord c of s_i that splits s_i into two disconnected components, each containing points from Q. Since $E_c \cup E_l$ properly connects s_i, there must be an edge $e' \in E_c \cup E_l$, whose endpoints belong to different components of $s_i \setminus c$. This implies that e' crosses c an odd number of times. Since e is an edge in the plane graph G_{l+1}, e' cannot be an edge in G_{l+1}. Therefore, $e' \in E_l \setminus E'_l$, meaning that e' pierces s_{l+1}. All the intersections between e' and c happen at segments of e' which are chords of s_{l+1}, since we excluded the interiors of the obstacles from I_{l+1}. Hence, there is one segment c' of e' which is a chord of s_{l+1} and has an odd number of intersections with c. Moreover, c' lies in the interior of s_i since it is contained in e', and similarly c lies in the interior of s_{l+1}. This contradicts Lemma 4 and therefore e cannot be piercing s_i.

Therefore, none of the edges in $E_c \cup E'_l \cup D'$ pierce any of $s_1, s_2, \cdots, s_l$, and moreover $E_c \cup E'_l \cup D'$ properly connects s_{l+1}. Using the induction hypothesis for regions $s_1, \cdots, s_l$ with $E_c \cup E'_l \cup D$ as the new compulsory set of edges, we obtain a plane multigraph whose edge set properly connects $s_1, s_2, \cdots, s_{l+1}$. ☐

Lemma 6. *Given a family S of k-admissible regions and a set of points P in the plane, there is a planar graph $G_Q = (Q, E_Q)$ for any $Q \subseteq P$, such that for any $s \in S$, $Q(s)$ induces a connected subgraph.*

Proof. The required planar graph is obtained by applying Lemma 5 for the family S and the set Q, with $E_c = \emptyset$ as the compulsory set of edges and replacing multi-edges with single edges in the resulting plane multigraph. ☐

Remark 2. *Since planar graphs are 4-colorable [5], the above Lemma implies that the primal hypergraph defined by a set of points and a set of k-admissible regions in the plane is 4-colorable. (The colorability of* dual *hypergraphs induced by geometric objects has been studied in [22].)*

Lemma 6 and Theorem 4 imply:

Theorem 11. *The dual hypergraph $\mathcal{H}^*(P, S)$ defined by a family S of k-admissible regions and a set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Showing the existence of a $o(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ size ϵ-net for the primal hypergraph $\mathcal{H}(P, S)$ is an open question. However, we prove the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ when S is a family of 2-admissible regions (pseudo-disks).

Lemma 7. *Let S be a family of pseudo-disks and P a set of points in $\mathbb{R}^2$. Then, for any $R \subseteq S$ there is a graph $G_R = (R, E_R)$, such that $|E_R| < 24|R|$, and such that for any $p \in P$, if p is contained in d pseudo-disks from R, then there are at least $\frac{1}{4}d - 1$ edges among those pseudo-disks.*

Proof. For a point $p \in P$, we denote by $R(p) \subseteq R$ the set of pseudo-disks in R that contain p and we define the *degree* $d(p)$ of p as $d(p) = |R(p)|$.

We start by constructing a set $P' \subseteq P$ as follows: Let $p_1, p_2, \cdots, p_n$ be all the points of P in decreasing order of their degrees. Insert p_i into P', if and only if, for every $p_j \in P'$ with $j < i$, we have $|R(p_i) \cap R(p_j)| < \frac{1}{4}d(p_i)$.

Using Lemma 6 for P' and R, we get a planar graph $G' = (P', E')$ such that for each $s \in R$, the points in $P'(s) = \{p \in P' : p \in s\}$ induce a connected subgraph. Therefore, for any region $s \in R$, it holds that $m_s \geq n_s - 1$, where n_s is the number of points of P' contained in s and m_s is the number of edges among those points. Summing over all $s \in R$, we get:

$$\sum_{s \in R} n_s \leq \sum_{s \in R} m_s + |R|. \tag{8}$$

Since for any $p_1, p_2 \in P'$, the sets $R(p_1)$ and $R(p_2)$ share fewer than $\frac{1}{4}\min\{d(p_1), d(p_2)\}$ pseudo-disks, we have:

$$\sum_{s\in R} m_s < \sum_{(p_1,p_2)\in E'} \frac{1}{4}\min\{d(p_1), d(p_2)\}.$$

Combining this with the fact that $\sum_{s\in R} n_s = \sum_{p\in P'} d(p)$, Inequality (8) gives

$$\begin{aligned}\sum_{p\in P'} d(p) &< \frac{1}{4}\sum_{(p_1,p_2)\in E'} \min\{d(p_1), d(p_2)\} + |R| \\ &\leq \frac{3}{4}\sum_{p\in P'} d(p) + |R|,\end{aligned}$$

where for the last inequality we are using the fact that since G' is planar, there is a way to orient the edges in E' so that the out-degree of every node in P' is at most 3 (see [8]). By rearranging, we get:

$$\sum_{p\in P'} d(p) < 4|R|. \tag{9}$$

Now we will construct G_R in such a way that for every $p \in P'$, $R(p)$ induces a connected subgraph. Consider some $p \in P'$ and the family $R(p)$ of pseudo-disks that contain it. Using Lemma 2.11 of [1], we obtain a combinatorially equivalent family of pseudo-disks that are all star-shaped with respect to p. From this, by performing an angular sweep using a ray emanating from p, (and having initial orientation such that it doesn't pass through any of the intersection points among the pseudo-disks) we get a combinatorially equivalent family of pseudo-parabolas. Applying Lemma 3 gives a graph $G_p = (R(p), E_p)$, with $|E_p| \leq 6|R(p)| = 6d(p)$, while for any $q \in P$, the pseudo-disks in $R(p) \cap R(q)$ induce a connected subgraph of G_p.

The union of the graphs G_p, for all $p \in P'$ gives the graph $G_R = (R, \cup_{p\in P'} E_p)$. Due to the way that P' was constructed, for any $p \in P \setminus P'$ there is some $p' \in P'$, such that $|R(p') \cap R(p)| \geq \frac{1}{4}d(p)$. Therefore, since the pseudo-disks of $R(p') \cap R(p)$ induce a connected subgraph in $G_{p'}$, there are at least $\frac{1}{4}d(p) - 1$ edges among the pseudo-disks that contain p. On the other hand, if $p \in P'$, then there are at least $d(p) - 1$ edges among the pseudo-disks in $R(p)$. We conclude the proof by observing that the total number of edges in G_R is at most $6(\sum_{p\in P'} d(p)) < 24|R|$ (using (9)). ☐

From Lemma 7 and Theorem 4 we get:

Theorem 12. *The primal hypergraph $\mathcal{H}(P, S)$ defined by a family S of pseudo-disks and a set P of points in the plane admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Remark 3. *Since pseudo-parabolic halfplanes are a special case of pseudo-disks, Theorem 12 implies that the primal hypergraph defined by a family of pseudo-parabolic halfplanes and a set of points in the plane, also admits an ϵ-net of size $O(\frac{1}{\epsilon})$.*

Remark 4. *Note that the definition of pseudo-disks that we used here is different than the one used in [18] where the family of pseudo-disks is required to be such that there is exactly one pseudo-disk passing any three non-collinear points. We do not make such an assumption and hence our result is stronger than the one in [18]. Moreover, Theorem 12 cannot be proved using the framework developed in [10], since their technique is applicaple only to* dual hypergraphs *(in the sense in which we have used the term in this paper, see Section 3) induced by geometric objects.*

7 Conclusions

We described a technique that reduces the problem of constructing small ϵ-nets for a hypergraph $\mathcal{H} = (X, S)$, to constructing a *small* graph on a subset $S' \subseteq S$ of the hyperedges so that the subset $S'(x) \subseteq S'$ of hyperedges containing a vertex x have $\Omega(|S'(x)|)$ edges among them. In the applications we considered the required graphs were obtained by easy constructive methods. However, the technique can be easily extended to allow the edges of the graph to have weights. In this case, we require that the total weight of the graph is *small* and at the same time the weight of the subgraph induced by $S'(x)$ is $\Omega(|S'(x)|)$. Such extensions may allow probabilistic techniques to be used for showing the existence of the required graph. As mentioned before, we are still far from properly characterizing hypergraphs that admit ϵ-nets of size $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$. It remains open whether hypergraphs induced by halfspaces and points in higher dimensions (> 3) also admit ϵ-nets of size $o(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$.

8 Acknowledgements

The authors wish to thank Rajiv Raman and Deepak Ajwani for their critical comments, and the anonymous referees for their helpful suggestions. The second author would like to thank Raimund Seidel, Hans Raj Tiwary and Nabil H. Mustafa for enthusiastic and non-trivial discussions.

References

[1] Pankaj K. Agarwal, Eran Nevo, János Pach, Rom Pinchasi, Micha Sharir, and Shakhar Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *J. ACM*, 51(2):139–186, 2004.

[2] Pankaj K. Agarwal, János Pach, and Micha Sharir. State of the union, of geometric objects: A review. manuscript, 2007.

[3] Pankaj K. Agarwal and Micha Sharir. Pseudo-line arrangements: duality, algorithms, and applications. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 800–809, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[4] Martin Anthony and Norman L. Biggs. *Computational Learning Theory: An Introduction.* Cambridge University Press, UK, 1992.

[5] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable. *Journal of Mathematics*, 21:108–121, 1977.

[6] Claude Berge. Hypergraphs. page 256, 1989.

[7] Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[8] Marek Chrobak and David Eppstein. Planar orientations with low out-degree and compaction of adjacency matrices. *Theor. Comput. Sci.*, 86(2):243–266, 1991.

[9] Kenneth L. Clarkson. Randomized geometric algorithms. *Computers and Euclidean Geometry*, 1992.

[10] Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.

[11] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational geometry: algorithms and applications.* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1997.

[12] David Haussler and Emo Welzl. ϵ-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[13] János Komlós, János Pach, and Gerhard J. Woeginger. Almost tight bounds for ϵ-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.

[14] Sören Laue. Geometric set cover and hitting sets for polytopes in R^3. In Susanne Albers and Pascal Weil, editors, *25th International Symposium on Theoretical Aspects of Computer Science (STACS 2008)*, Dagstuhl, Germany, 2008. Internationales Begegnungs- und Forschungszentrum f"ur Informatik (IBFI), Schloss Dagstuhl, Germany.

[15] Jirí Matousek. Reporting points in halfspaces. *Comput. Geom.*, 2:169–186, 1992.

[16] Jirí Matousek. Epsilon-nets and computational geometry. *New Trends in Discrete and Computational Geometry, Algorithms and Combinatorics*, 10:69–89, 1993.

[17] Jirí Matousek. *Lectures in Discrete Geometry.* Springer-Verlag, New York, NY, 2002.

[18] Jirí Matousek, Raimund Seidel, and Emo Welz. How to net a lot with little: Small ϵ-nets for disks and halfspaces. In *Symposium on Computational Geometry*, pages 16–22, 1990.

[19] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms.* Cambridge University Press, New York, 1995.

[20] Ketan Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms.* Prentice Hall, Englewood Cliffs, NJ, 1994.

[21] János Pach and Gerhard J. Woeginger. Some new bounds for epsilon-nets. In *Symposium on Computational Geometry*, pages 10–15, 1990.

[22] Shakhar Smorodinsky. On the chromatic number of some geometric hypergraphs. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 316–323, New York, NY, USA, 2006. ACM.

[23] Vladimir N. Vapnik and Alexey Ya. Chervonenkis. On the uniform convergence of relative frequencies to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, 1971.