The ϵ-t-Net Problem

Noga Alon
Department of Mathematics, Princeton University, Princeton, NJ 08544, USA
Schools of Mathematics and Computer Science, Tel Aviv University, Tel Aviv 69978, Israel
nogaa@tau.ac.il

Bruno Jartoux
Department of Computer Science, Ben-Gurion University of the Negev, Be'er-Sheva, Israel
jartoux@post.bgu.ac.il

Chaya Keller
Department of Computer Science, Ariel University, Ariel, Israel
chayak@ariel.ac.il

Shakhar Smorodinsky
Department of Mathematics, Ben-Gurion University of the Negev, Be'er-Sheva, Israel
shakhar@math.bgu.ac.il

Yelena Yuditsky
Department of Mathematics, Ben-Gurion University of the Negev, Be'er-Sheva, Israel
yuditskyL@gmail.com

──── **Abstract** ────

We study a natural generalization of the classical ϵ-net problem (Haussler–Welzl 1987), which we call *the ϵ-t-net problem*: Given a hypergraph on n vertices and parameters t and $\epsilon \geq \frac{t}{n}$, find a minimum-sized family S of t-element subsets of vertices such that each hyperedge of size at least ϵn contains a set in S. When $t = 1$, this corresponds to the ϵ-net problem.

We prove that any sufficiently large hypergraph with VC-dimension d admits an ϵ-t-net of size $O(\frac{(1+\log t)d}{\epsilon} \log \frac{1}{\epsilon})$. For some families of geometrically-defined hypergraphs (such as the dual hypergraph of regions with linear union complexity), we prove the existence of $O(\frac{1}{\epsilon})$-sized ϵ-t-nets.

We also present an explicit construction of ϵ-t-nets (including ϵ-nets) for hypergraphs with bounded VC-dimension. In comparison to previous constructions for the special case of ϵ-nets (i.e., for $t = 1$), it does not rely on advanced derandomization techniques. To this end we introduce a variant of the notion of VC-dimension which is of independent interest.

1   Introduction

1.1   Preliminaries

Hypergraphs and VC-dimension.

A *hypergraph* is a pair $H = (V, \mathcal{E})$ where V is a set of *vertices* and $\mathcal{E} \subseteq 2^V$ is the set of *hyperedges* of H. When V is finite, H is a *finite* hypergraph.

A subset $V' \subseteq V$ is *shattered* if all its subsets are realized by $\mathcal{E}$, meaning $\{V' \cap e : e \in \mathcal{E}\} = 2^{V'}$. The *VC-dimension* of H, denoted by $\dim H$, is the cardinality of a largest shattered subset of V or $+\infty$ if arbitrarily large subsets are shattered (which does not happen in finite hypergraphs). This parameter plays a central role in statistical learning, computational geometry, and other areas of computer science and combinatorics [44, 31, 33].

ϵ-nets, Mnets.

Let $\epsilon \in (0, 1)$. An ϵ-*net* for a finite hypergraph $(V, \mathcal{E})$ is a subset of vertices $S \subseteq V$ such that $S \cap e \neq \emptyset$ for every hyperedge $e \in \mathcal{E}$ such that $|e| \geq \epsilon|V|$.

Haussler and Welzl [22] proved that finite hypergraphs with VC-dimension d admit ϵ-nets of size $O(\frac{d}{\epsilon} \log \frac{d}{\epsilon})$, later improved to $O(\frac{d}{\epsilon} \log \frac{1}{\epsilon})$ [27]. In the last three decades, ϵ-nets have found applications in diverse areas of computer science, including machine learning [10], algorithms [14], computational geometry [7] and social choice [4].

Mustafa and Ray introduced the notion of *Mnets* [32]. For a hypergraph $(V, \mathcal{E})$ and for a fixed $\epsilon \in (0, 1)$, an ϵ-*Mnet* is a family $\{V_1, V_2, \ldots, V_\ell\}$ such that each $V_i \subseteq V$, each V_i is of size $\Theta(\epsilon|V|)$, and, for each $e \in \mathcal{E}$ such that $|e| \geq \epsilon|V|$, $V_i \subseteq e$ for some V_i. They constructed small ϵ-Mnets (i.e., such families with small ℓ) for several classes of geometric hypergraphs. These results were extended by Dutta et al. [19] using polynomial partitioning.

Explicit constructions

Although Hausssler and Welzl's proof of the ϵ-net theorem is probabilistic, several deterministic constructions of ϵ-nets for hypergraphs with finite VC-dimension have been devised [11, 29, 15]. The best result of this kind is Brönniman, Chazelle and Matoušek's $O(\epsilon^{-d} \log^d \frac{1}{\epsilon}|V|)$-time algorithm for computing an ϵ-net of size $O(\frac{d}{\epsilon} \log \frac{d}{\epsilon})$ [11]. These constructions are used to derandomize applications of ϵ-nets, such as low-dimensional linear programming [14].

In scenarios where the VC-dimension is $\Omega(\log|V|)$, the running time of these constructions becomes exponential in $|V|$. For one such scenario – the hypergraph induced by half-spaces on the discrete cube $V = \{-1, 1\}^d$ – Rabani and Shpilka [36] presented an efficient explicit construction of an ϵ-net, alas of sub-optimal size: $O(\epsilon^{-b}|V|^a)$ for some universal constants $a, b > 0$, whereas $O(|V|/\epsilon)$ can be obtained by random sampling. Like the aforementioned explicit constructions, the construction of [36] is based on derandomization.

1.2   Our problem

We denote by $\binom{X}{k}$ the set of all subsets of cardinality k (or "k-subsets") of the set X.

▶ **Definition 1.** *Let $H = (V, \mathcal{E})$ be a finite hypergraph, t a positive integer and $\epsilon \in (t/|V|, 1)$. A family $S \subseteq \binom{V}{t}$ of t-subsets of V is an ϵ-t-net for H if for every $e \in \mathcal{E}$ with $|e| \geq \epsilon|V|$ there is an $s \in S$ such that $s \subseteq e$.*

As mentioned already, for $t = 1$ this is equivalent to the ϵ-net notion, and for $t = \Theta(\epsilon|V|)$ this corresponds to the notion of ϵ-Mnets. In this paper we study the following problem.

▶ **Problem.** *How small are the smallest ϵ-t-nets for H? Can we compute them efficiently?*

Motivation.

Instances of the ϵ-t-net problem appear naturally in various contexts in computer science and combinatorics. For example, the following is a basic motivating example for *secret sharing* [28, 40]: "Eleven scientists are working on a secret project. They wish to lock up the documents in a cabinet so that the cabinet can be opened if and only if six or more of the scientists are present. What is the smallest number of locks needed?". Consider a variant of this question in which the number of scientists is large. We still insist on the basic security condition – that no less than six scientists can open the cabinet. On the other hand, due to the large number of scientists, we do not require that any six should be able to do so, but rather any sufficiently large group of a certain kind, e.g., at least one tenth of all scientists including a representative of each university involved.

The classical secret sharing methods (see, e.g., [9]) distribute "keys" to subsets of 6 scientists so that any six scientists will be able to open the cabinet but no five will be able to do that. But as we require only certain groups of scientists to be able to open it, it is possible to distribute shared keys to only some of the 6-subsets. The questions: "What is the minimal number of 6-subsets we can achieve? and how can we choose the 6-subsets of scientists we distribute keys to?" are an instance of the ϵ-t-net problem – with $t = 6$, $\epsilon = 1/10$, and the hyperedges of the hypergraph being all groups of scientists that are required to be able to open the cabinet.

Other contexts in which the ϵ-t-net problem appears (described in Section 6) include the Turán numbers of hypergraphs, χ-boundedness of graphs, edge-coloring of hypergraphs and more.

Related work: ϵ-Nets and Mnets.

For any t, the minimum size of an ϵ-t-net is sandwiched between the corresponding minimum sizes of ϵ-nets and of Mnets. Indeed, given an Mnet, one obtains an ϵ-t-net by picking one t-subset from each subset, and given an ϵ-t-net, one obtains an ϵ-net by taking one vertex from each t-subset. The survey [33] has most known bounds on these objects.

1.3   Results

Notation: we write $O_{x,y}(\cdot)$ when the implicit constants depend on parameters x and y.

Hypergraphs of finite VC-dimension have small ϵ-t-nets.

Our main result is an existence result for small ϵ-t-nets.

▶ **Theorem 2.** *For every $\epsilon \in (0,1)$ and $t \in \mathbb{N} \setminus \{0\}$, every hypergraph on $\geq C_1 \left(\frac{t-1}{\epsilon}\right)^{d^*}$ vertices with VC-dimension d and dual shatter function $\pi_H^*(m) \leq Cm^{d^*}$ admits an ϵ-t-net of size $O(\frac{d(1+\log t)}{\epsilon} \log \frac{1}{\epsilon})$, all elements of which are pairwise disjoint. Here $C_1 = C_1(d^*, C)$.*

(The dual shatter function, described in Section 2, is a property of the hypergraph such that we may always take $d^* < 2^{d+1}$.)

This bound is asymptotically tight when $t = O(1)$, in the sense that there exist hypergraphs for which any ϵ-net, and consequently also any ϵ-t-net, is of size $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ [27]. The proof of Theorem 2 involves a surprising relation between the ϵ-t-net problem and the existence of *spanning trees with a low crossing number*, proved by Welzl in 1988 [46].

Hypergraphs with VC-dimension 1 admit $O(\frac{1}{\epsilon})$-sized ϵ-nets [27] and ϵ-Mnets [19]. The latter fact yields the following result, albeit with worse constants. We offer a simple proof.

▶ **Theorem 3.** *For every positive integer t and $\epsilon \le \frac{1}{2}$, every finite hypergraph on $\ge t\lceil\frac{1}{\epsilon}\rceil$ vertices with VC-dimension 1 admits an ϵ-t-net of size at most $t\lceil\frac{1}{\epsilon}\rceil + 1$.*

An efficient explicit construction of ϵ-t-nets.

Our second result is a new explicit construction of ϵ-t-nets, for all $t \ge 1$. The case of $t = 1$ (i.e., ϵ-nets) is of independent interest, as in this case our construction does not follow the proof strategy of Haussler and Welzl and does not use derandomization (unlike all previously known explicit constructions of ϵ-nets). On the other hand, it has a sub-optimal size of $O_d(\frac{1}{\epsilon^d})$, where d is the VC-dimension of the underlying hypergraph.

For a higher t, we introduce a new parameter of the hypergraph, which we call the *t-VC-dimension*. For hypergraphs of t-VC-dimension d, we construct ϵ-t-nets of size $O_d(\frac{1}{\epsilon^{d+t-1}})$. We give some first results on the relation between this new parameter and the standard VC-dimension.

Small ϵ-2-nets for geometric hypergraphs.

In view of Theorem 2, which shows that for hypergraphs with a constant VC dimension one can obtain an ϵ-t-net of roughly the same size as the smallest ϵ-net, it is natural to ask whether a similar result can be achieved for geometrically-defined hypergraphs that admit an ϵ-net of size $O(\frac{1}{\epsilon})$. We obtain such results for several geometrically-defined hypergraphs in $\mathbb{R}^2$, including the intersection hypergraph of two families of pseudo-disks and the dual hypergraph of a family of regions with linear union complexity. Namely, we show that these hypergraphs have $O(\frac{1}{\epsilon})$-sized ϵ-2-nets provided they have $\Omega(\frac{1}{\epsilon})$ vertices. Interestingly, in some scenarios the minimum size of an ϵ-2-net is sensitive to the exact multiplicative constant: there are subhypergraphs on $\Theta(\frac{1}{\epsilon})$ vertices for which any ϵ-2-net is of size $\Omega(\frac{1}{\epsilon^2})$.

2   Construction of Auxiliary Hypergraphs

2.1   Some preparatory results

Sauer's lemma.

Given a hypergraph $H = (V, \mathcal{E})$ the *trace (also known as projection or restriction) of H on* $A \subseteq V$ is $\Pi_H(A) = \{A \cap e \colon e \in \mathcal{E}\}$; shattered subsets are those for which $\Pi_H(A) = 2^A$. The shatter function of H is

$$\pi_H \colon n \in \mathbb{N} \mapsto \max\{|\Pi_H(A)| : A \subseteq V, \ |A| \le n\}.$$

It is bounded by the *Sauer–Shelah lemma*:

▶ **Lemma 4** ([44, 38, 41]). *If $\dim H = d$ then $\pi_H(n) \le \binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{d}$. In particular, for $1 \le d \le n$ one has $\pi_H(n) \le (\frac{e}{d})^d \cdot n^d$, where e is Euler's number.*
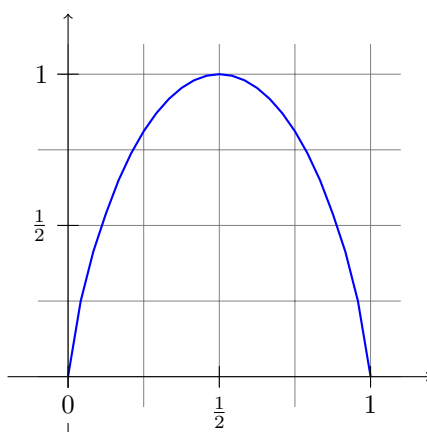
Figure 1 The binary entropy function.

Binary entropy function.

This is $h\colon x \in (0,1) \mapsto -x \log x - (1-x) \log(1-x)$. (All logarithms are binary. See Figure 1.) We will use the following inequality.

$$\forall \alpha \in \left(0, \frac{1}{2}\right], \ \forall n \in \mathbb{N}, \quad \log \sum_{i=0}^{\lfloor \alpha n \rfloor} \binom{n}{i} \leq nh(\alpha). \tag{1}$$

Proof. For the sake of simplicity, we assume $\alpha n \in \mathbb{N}$. By the binomial theorem,

$$\begin{aligned}
1 = (\alpha + (1-\alpha))^n &= \sum_{i=0}^{n} \binom{n}{i} \alpha^i (1-\alpha)^{n-i} \\
&\geq \sum_{i=0}^{\alpha n} \binom{n}{i} \alpha^i (1-\alpha)^{n-i} = \sum_{i=0}^{\alpha n} \binom{n}{i} (1-\alpha)^n \left(\frac{\alpha}{1-\alpha}\right)^i \\
&\geq \sum_{i=0}^{\alpha n} \binom{n}{i} (1-\alpha)^n \left(\frac{\alpha}{1-\alpha}\right)^{\alpha n} \quad \text{since } 0 < \frac{\alpha}{1-\alpha} < 1 \\
&= 2^{-n \cdot h(\alpha)} \cdot \sum_{i=0}^{\alpha n} \binom{n}{i}. \qquad\qquad\qquad\qquad\qquad\qquad \blacktriangleleft
\end{aligned}$$

The binary entropy function restricted to $(0, \frac{1}{2}]$ is invertible, and [13, Th. 2.2]:

$$\forall x \in (0,1), \quad \frac{x}{2\log \frac{6}{x}} \leq h^{-1}(x) \leq \frac{x}{\log \frac{1}{x}}. \tag{2}$$

2.2 A first hypergraph on t-subsets

▶ **Definition 5.** *Given a hypergraph $H = (V, \mathcal{E})$ and a positive integer t, let H^t be the hypergraph $(V^t, \mathcal{E}^t)$ where $V^t = \binom{V}{t}$ and $\mathcal{E}^t = \left\{\binom{e}{t}\colon e \in \mathcal{E}\right\}$. That is, its vertices are all t-element subsets of V and each hyperedge of H^t consists of all such subsets contained in a given hyperedge of H.*

For $t \in \mathbb{N} \setminus \{0, 1\}$, let $\gamma_t = (th^{-1}(1/t))^{-1}$. Note that $\log t \leq \gamma_t \leq 2\log 6t$.

▶ **Proposition 6.** *If H is a hypergraph with $\dim H = d$ then $d - t + 1 \leq \dim H^t \leq \gamma_t d$.*

Proof. We assume that $t \geq 2$, as for $t = 1$, $\dim H^t = \dim H^1 = \dim H = d$.

To prove the left inequality, let $\{v_1, \ldots, v_d\}$ be a shattered subset of vertices in H, with $d \geq t - 1$. There are $d - t + 1$ sets containing all vertices in $\{v_1, v_2, \ldots, v_{t-1}\}$ and exactly one in $\{v_t, v_{t+1}, \ldots v_d\}$. It is easy to see that they form a shattered subset in H^t.

For the right inequality, suppose to the contrary that P is a shattered set in H^t with $d' = |P| > \gamma_t d$. Let $S = \cup_{p \in P} p$; clearly $|S| \leq td'$. Observe also that $d' + t - 1 \leq |S|$. If this were not the case there would exist some $p_1 \in P$ such that $p_1 \subseteq \cup_{p \in P \setminus \{p_1\}} p$, which would contradict the fact that P is shattered.

We denote $|S| = \beta d'$; we have $1 < \beta \leq t$.

Since P is shattered in H^t, each $P_1 \subseteq P$ is of the form $P \cap \binom{e}{t} = \{p \in P : p \subseteq (S \cap e)\}$ for some $e \in \mathcal{E}$. Thus $|\Pi_H(S)| \geq |2^P| = 2^{d'}$.

On the other hand, $\dim H = d$, and so by Lemma 4, $|\Pi_H(S)| \leq \binom{\beta d'}{0} + \binom{\beta d'}{1} + \cdots + \binom{\beta d'}{d}$. It follows from Equation (1) (with $\beta d' \geq \beta \gamma_t d > 2d$) that $d' \leq \log|\Pi_H(S)| \leq \beta d' h(\frac{d}{\beta d'})$. We show that $1 > \beta h(\frac{d}{\beta d'})$, a contradiction.

Note that $\frac{1}{t\gamma_t} \leq \frac{1}{\gamma_t \beta} < \frac{1}{2}$. Since $t \mapsto \frac{h(t)}{t}$ is monotone decreasing in the range $(0, 1)$, we have $\gamma_t \beta \cdot h(\frac{1}{\gamma_t \beta}) \leq t\gamma_t \cdot h(\frac{1}{t\gamma_t}) = \gamma_t$. As h is increasing on $(0, \frac{1}{2})$, it follows that $\beta h(\frac{d}{\beta d'}) < \beta h(\frac{1}{\beta \gamma_t}) \leq 1$. ◄

Proposition 6 allows us to slightly improve the "trivial" upper bound of $O(\frac{d^t}{\epsilon^t}(\log \frac{1}{\epsilon})^t)$ on the minimum size of an ϵ-t-net for any hypergraph with constant VC-dimension.

▶ **Corollary 7.** *Let H be a hypergraph on n vertices with VC-dimension d. For any t, ϵ such that $n \geq \frac{t}{\epsilon}$, H admits an ϵ-t-net of size $O(\frac{dt(1+\log t)}{\epsilon^t} \log \frac{1}{\epsilon})$.*

Indeed, observe that an ϵ^t-net for H^t is an ϵ-t-net for H, and apply the classical ϵ-net theorem to H^t.

2.3 A smaller, well-behaved hypergraph on t-subsets

A *spanning cycle* P for $H = (V, \mathcal{E})$ is a cycle graph on V that visits all vertices (exactly once). For $e \in \mathcal{E}$, let $\mathrm{cr}(P, e)$ be the number of edges of P with one endpoint in e and the other in $V \setminus e$. The *crossing number* of P with respect to H is $\sup\{\mathrm{cr}(P, e) : e \in \mathcal{E}\}$.

The *dual hypergraph* of H is $H^* = (\mathcal{E}, \mathcal{E}^*)$, where $\mathcal{E}^*$ consists of all hyperedges $v^* = \{e \in \mathcal{E} : v \in e\}$ for $v \in V$. Its shatter function is the *dual shatter function* of H, and is denoted by π_H^*.

If $\dim H = d$ then $\dim H^* \leq 2^{d+1}$ [8], and hence $\pi_H^*(m) \leq C_d m^{2^{d+1}}$ for every positive m, where C_d is a constant depending on d. In particular, any hypergraph with finite VC-dimension satisfies the hypotheses of the following theorem.

▶ **Theorem 8** ([46, Lemma 3.3 and Theorem 4.2])**.** *Let H be a hypergraph on n vertices such that $\pi_H^*(m) \leq Cm^d$ for some constants $C > 0$ and $d > 1$. Then there exists another constant C_1 (depending on C and d) and a spanning cycle for H with crossing number $\leq C_1 n^{1-\frac{1}{d}}$.*

(An additional $\log n$ factor in Welzl's original result was later removed [30, Sec. 5.4]. Up to constant factors, this theorem is equivalent to the same result for paths or trees.)

▶ **Definition 9.** *Let $H = (V, \mathcal{E})$ be a finite hypergraph with $\pi_H^*(m) \leq Cm^d$. Let P be a spanning cycle for H whose crossing number is minimal (and thus $\leq C_1|V|^{1-\frac{1}{d}}$). Fix an arbitrary starting point $v_0 \in P$ and orientation of P. For $0 \leq i < |V|$, let $v_i \in V$ be the i-th vertex along P. Let $V_{lc}^t = \{\{v_{kt}, v_{kt+1}, \ldots, v_{kt+t-1}\} : 0 \leq k < \lfloor\frac{|V|}{t}\rfloor\} \subsetneq \binom{V}{t}$ (where the subscript lc stands for low crossing). Observe that its elements are pairwise disjoint. Let H_{lc}^t be the hypergraph on V_{lc}^t whose hyperedges are of the form $\{v \in V_{lc}^t : v \subseteq e\}$ for each $e \in \mathcal{E}$.*

▶ **Remark 10.** In order to make H_{lc}^t uniquely defined, P is chosen arbitrarily from all suitable spanning cycles. As H_{lc}^t is a subhypergraph of H^t, $\dim H_{lc}^t \leq \dim H^t$, and thus we also have $\dim H_{lc}^t \leq \gamma_t \dim H$.

3   Existence of Small ϵ-t-Nets

▶ **Theorem 2.** *For every $\epsilon \in (0,1)$ and $t \in \mathbb{N} \setminus \{0\}$, every hypergraph on $\geq C_1 \left(\frac{t-1}{\epsilon}\right)^{d^*}$ vertices with VC-dimension d and dual shatter function $\pi_H^*(m) \leq Cm^{d^*}$ admits an ϵ-t-net of size $O(\frac{d(1+\log t)}{\epsilon} \log \frac{1}{\epsilon})$, all elements of which are pairwise disjoint. Here $C_1 = C_1(d^*, C)$.*

Proof. For $t = 1$, this is simply the ϵ-net theorem. For higher t, let $H = (V, \mathcal{E})$ be such a hypergraph and $n = |V|$. Consider the hypergraph H_{lc}^t defined in Section 2. It has $\lfloor \frac{n}{t} \rfloor$ vertices and VC-dimension $\leq \gamma_t d$ (by Remark 10), and thus admits an $\frac{\epsilon}{2}$-net of size $O(\frac{\gamma_t d}{\epsilon} \log \frac{1}{\epsilon})$. We claim that any such $\frac{\epsilon}{2}$-net $N \subseteq \binom{V}{t}$ is also an ϵ-t-net for H.

Indeed, the crossing number of the associated spanning cycle is $O_{C,d^*}(n^{1-1/d^*})$. Every hyperedge e of H with $|e| \geq \epsilon n$ fully contains at least $\lfloor \frac{\epsilon n}{t} \rfloor - O_{C,d^*}(n^{1-1/d^*})$ elements of V_{lc}^t, which is $\geq \frac{\epsilon n}{2t}$ as soon as $n = \Omega_{C,d^*}(\frac{t}{\epsilon} n^{1-1/d^*})$, or equivalently (noting also that $2(t-1) \geq t$ for $t \geq 2$) when $n = \Omega_{C,d^*}(\frac{t-1}{\epsilon}^{d^*})$. One of these t-subsets is in N.   ◄

▶ **Remark 11.** In general, some fast growth of $n = |V|$ as a function of $\frac{1}{\epsilon}$ is necessary. For example, given any ϵ such that $\frac{t}{\epsilon} \in \mathbb{N}$, the complete t-uniform hypergraph on $\frac{t}{\epsilon}$ vertices does not have any ϵ-t-net with fewer than $\binom{t/\epsilon}{t}$ elements. Moreover, there exist geometrically-defined hypergraphs that do not admit ϵ-2-nets of size $o(\frac{1}{\epsilon^2})$ (see Figure 2 and subsection 5.3.2). On the other hand, in Section 5 we show that certain classes of geometrically-defined hypergraphs have "small" ϵ-t-nets even for "small" values of n.

Small ϵ-nets, small ϵ-2-nets.

A natural question arising from Theorem 2 is whether any hypergraph that admits small ϵ-nets must also admit ϵ-t-nets of approximately same size. In general, the answer is negative. Take for example a hypergraph whose smallest ϵ-net is of size $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ (see [27], [34]), and augment it by adding a vertex that belongs to all hyperedges. Clearly, this second hypergraph has the same VC-dimension and a one-element ϵ-net, but any ϵ-2-net is of size $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$.

However, this example is quite artificial. In "natural" scenarios (and for sufficiently large vertex sets) the smallest ϵ-nets and ϵ-2-nets might still have approximately same size. In Section 5 we show that this is the case for some geometrically-defined hypergraphs.

Another scenario in which there exist both an ϵ-net and an ϵ-2-net of size $O(\frac{1}{\epsilon})$ is when the VC-dimension of the hypergraph is 1. In this case, the existence of an ϵ-net of size $O(\frac{1}{\epsilon})$ was proved in [27]. The next theorem could be derived from results on Mnets [19], at the cost of poor multiplicative constants. Here we give a simpler proof for it.

▶ **Theorem 3.** *For every positive integer t and $\epsilon \leq \frac{1}{2}$, every finite hypergraph on $\geq t\lceil \frac{1}{\epsilon} \rceil$ vertices with VC-dimension 1 admits an ϵ-t-net of size at most $t\lceil \frac{1}{\epsilon} \rceil + 1$.*

Proof. Let $(V, \mathcal{E})$ be such a hypergraph and $n = |V|$. Without loss of generality, $\min\{|e| : e \in \mathcal{E}\} \geq \epsilon n$. For $1 \leq i \leq t$, there exists an ϵ-net N_i that hits each $e \in \mathcal{E}$ at least i times, and $|N_i| = i\lceil \frac{1}{\epsilon} \rceil$. To see this let N_1 be an ϵ-net of size $\lceil \frac{1}{\epsilon} \rceil$ [27]. In the hypergraph induced on $V \setminus N_i$ the hyperedges hit only i times by N_i have cardinality $\geq \epsilon n - i$, while the number of vertices is $n - i\lceil \frac{1}{\epsilon} \rceil$, for a ratio $\frac{\epsilon n - i}{n - i\lceil \frac{1}{\epsilon} \rceil} \geq \epsilon$. Take an ϵ-net N of size $\lceil \frac{1}{\epsilon} \rceil$ for this hypergraph and let $N_{i+1} = N_i \cup N$. Finally, let the desired ϵ-t-net consist of one t-subset from each element of $\Pi_H(N_t)$ with $\geq t$ vertices, of which there are at most $|N_t| + 1$ by Lemma 4.   ◄

4   Deterministic Construction of ϵ-t-Nets

Let $H = (V, \mathcal{E})$ be a finite hypergraph with VC-dimension d, and fix $\epsilon \in (0, 1)$. In this section we provide an explicit polynomial-time construction of ϵ-nets that immediately implies an explicit construction of ϵ-t-nets. The size is far from optimal, but the construction is simpler than previous explicit constructions, as it does not rely on packing numbers nor on pseudo-random choices.

4.1   Deterministic construction of ϵ-nets

We start with the following definition:

▶ **Definition 12.** *Let A, B be two subsets of V. We say that A stabs B if for every hyperedge $S \in \mathcal{E}$ with $B \subseteq S$ we have $S \cap A \neq \emptyset$.*

Let $S \in \mathcal{E}$ be a hyperedge, $|S| \geq d+1$, and let $X \in \binom{S}{d+1}$. Since the VC-dimension is d the set X is not shattered. Notice that $X = X \cap S \in \Pi_H(X)$. We can also assume that $\emptyset \in \Pi_H(X)$, for otherwise X is a transversal for H of size $d+1$. Hence there exists at least one non-trivial, proper subset $A \subsetneq X$ such that $(X \setminus A) \notin \Pi_H(X)$. Equivalently, there is a non-trivial partition of X into A and $X \setminus A$ such that A stabs $X \setminus A$. We say that X is of type $|A| \in \{1, \ldots, d\}$. Note that X could have several types. By the pigeonhole principle, there is a type i and a subset $A \in \binom{S}{i}$ such that a fraction $d^{-1}\binom{|S|}{i}^{-1}$ of the elements of $\binom{S}{d+1}$ are stabbed by A, hence the following lemma holds:

▶ **Lemma 13.** *Let S be a hyperedge containing $\geq d+1$ vertices of V. Then there exists an integer $i \in \{1, \ldots, d\}$ and a subset $A \in \binom{S}{i}$ that stabs $\binom{|S|}{d+1} d^{-1} \binom{|S|}{i}^{-1}$ subsets of cardinality $d+1-i$.*

Constructing ϵ-nets.

Put $n := |V|$. We construct an ϵ-net of size $O_d(\frac{1}{\epsilon^d})$ as follows. Start with $N = \emptyset$. As long as there is a hyperedge $S \in \mathcal{E}$ with $|S| \geq \epsilon n$ and $S \cap N = \emptyset$, Lemma 13 asserts that some i-subset from S stabs $\Omega_d((\epsilon n)^{d+1-i})$ subsets of S with cardinality $d+1-i$ for an appropriate $i \in \{1, \ldots, d\}$. Add all elements of this subset to N; we call this a type i iteration.

The resulting set is an ϵ-net by construction. It is left to show that $|N| = O_d(\frac{1}{\epsilon^d})$. As each step of the construction adds at most d vertices to N it is enough to bound the number of iterations T. By the pigeonhole principle, at least $\frac{T}{d}$ of the iterations have the same type, say i. After a type-i iteration N stabs an additional $\Omega_d((\epsilon n)^{d+1-i})$ subsets of cardinality $d+1-i$ none of which were previously stabbed. Since there are $\binom{n}{d+1-i}$ subsets of cardinality $d+1-i$ we have $\frac{T}{d} = O_d(\binom{n}{d+1-i}(\epsilon n)^{-(d+1-i)}) = O_d(\frac{1}{\epsilon^d})$.

Complexity analysis

We analyze the running time of the above algorithm. We assume that for the algorithm we have a data structure which is the incidence matrix of the hypergraph H. Without loss of generality, each hyperedge of $\mathcal{E}$ may be replaced with a subset of cardinality $\lceil \epsilon n \rceil$. This can be done in time $O(\epsilon n^{d+1})$ due to the fact that $|\mathcal{E}| = O(n^d)$.

We consider each $X \in \binom{V}{d+1}$. Firstly we check if there is a hyperedge $S \in \mathcal{E}$ which contains X, if not, we continue to the next subset. If yes, we consider each of the $2^{d+1} - 2$ proper subsets of X. Let $A \subset X$ be such a subset. We check if $X \setminus A$ is stabbed by A. We can do it by going over all $O(n^d)$ hyperedges of H. Hence, in total this pre-processing step

takes $O(n^{d+1} \cdot 2^{d+1} \cdot n^d) = O_d(n^{2d+1})$ running time. While determining the type of any $(d+1)$-subset of X and scanning all the hyperedges of the hypergraph, we maintain for any i-subset $A \subset X$ $(1 \leq i \leq d)$, a list of all the $(d+1-i)$-subsets of X that A stabs and their number.

Consider some iteration of the algorithm and let $S \in \mathcal{E}$ be such that $|S| \geq \epsilon n$ and $S \cap N = \emptyset$ where N is the collection of elements found until this iteration. We find a subset $A \subset S$ of size at most d which stabs the most subsets of size $(d+1) - |A|$.

The running time of each iteration is $O(|S|^d \cdot n^d) = O(\epsilon^d n^{2d})$. Hence in total the running time of the algorithm after the pre-processing step is $O_d(\frac{\epsilon^d n^{2d}}{\epsilon^d}) = O_d(n^{2d})$. Hence the total running of the algorithm described in the previous section is $O_d(n^{2d})$.

Immediate applications to ϵ-t-nets

The construction of ϵ-nets in Section 4.1 gives two straightforward constructions of ϵ-t-nets.
1. *Trivial construction.* Use the above algorithm to explicitly construct t disjoint ϵ-nets of size $O_d(1/\epsilon^d)$, and take all t-subsets of elements in their union that contain one element from each net. The resulting ϵ-t-net is of size $O_d(1/\epsilon^{td})$.
2. *Construction via H_{lc}^t.* Use the above algorithm to explicitly construct an $\frac{\epsilon}{2}$-net for the hypergraph H_{lc}^t, which is an ϵ-t-net for H (as was shown in the proof of Theorem 2). The resulting ϵ-t-net is of size $O_{d,t}(1/\epsilon^{\dim H_{lc}^t})$. (The cycle with a low crossing number required for constructing the hypergraph H_{lc}^t can be found in polynomial time [46, 30]).

4.2 Deterministic construction of ϵ-t-nets

We present a direct construction of ϵ-t-nets without passing through ϵ-nets. For the sake of convenience, we start by for presenting the method for $t = 2$.

The following definition extends the classical notion of VC-dimension.

▶ **Definition 14.** *Let t be a positive integer. Also let $H = (V, \mathcal{E})$ be a hypergraph, and T', T such that $T' \subseteq T \subseteq V$. We say that T' is t-realized by H (with respect to T) if $T' \cup S \in \Pi_H(T)$ for some $S \subseteq T$ such that $|S| < t$. We say that T is t-shattered by H if every $T' \subseteq T$ is t-realized by H (with respect to T). The t-VC-dimension of H, denoted by $\dim_t H$, is the maximal size of a vertex set that is t-shattered by H.*

Note that the 1-VC-dimension is the standard VC-dimension. Moreover, the t-VC-dimension is at most the $(t+1)$-VC-dimension for any positive integer t. We use the following adaptation of Definition 12:

▶ **Definition 15.** *Let $H = (V, \mathcal{E})$ be a hypergraph. Given two vertex sets $A, B \subseteq V$, we say that A 2-stabs B if each hyperedge of $\mathcal{E}$ that contains B also contains at least two vertices from A.*

▶ **Theorem 16.** *For a hypergraph $H = (V, \mathcal{E})$ with 2-VC-dimension d, one can construct explicitly an ϵ-2-net of size $O_d(1/\epsilon^{d-1})$.*

Proof. Let $S \in \mathcal{E}$ be a hyperedge and let $X \in \binom{S}{d+1}$. Since the 2-VC-dimension is d the set X is not 2-shattered. Notice that $X = X \cap S$ and so X and all elements of $\binom{X}{d}$ are 2-realized by H with respect to X. For our purpose, we can also assume that $\emptyset$ is 2-realized by H (with respect to X), for otherwise $\binom{X}{2}$ is a transversal for H of size $\binom{d+1}{2}$. This means that there is a partition, say $X = A \cup (X \setminus A)$, such that A 2-stabs $X \setminus A$. Let $i = |A|$. Note that $i \in \{2, \ldots, d\}$. We say that $X = A \cup (X \setminus A)$ is a type i partition. We need the following lemma, whose proof is similar to that of Lemma 13.

▶ **Lemma 17.** *Let S be a hyperedge containing $\geq d + 1$ vertices of V. Then there exists an integer $i \in \{2, \ldots, d\}$ and a subset $A \subset S$ with cardinality i that 2-stabs $\frac{\binom{|S|}{d+1}}{(d-1)\binom{|S|}{i}}$ subsets B of cardinality $d + 1 - i$.*

Constructing ϵ-2-nets

Let $H = (V, \mathcal{E})$ be as above and let $\epsilon > 0$ be fixed. Put $n = |V|$. We construct an ϵ-2-net of size $O_d(\frac{1}{\epsilon^{d-1}})$ as follows. We start with a set $N = \emptyset$. As long as there is a hyperedge $S \in \mathcal{E}$ with $|S| \geq \epsilon n$ that does not contain any pair $\{v, w\} \in N$, for an appropriate $i \in \{2, \ldots, d\}$ we take an i-subset $A \subset S$ 2-stabbing $\Omega_d((\epsilon n)^{d+1-i})$ subsets of S with cardinality $d + 1 - i$, and add to N all $\binom{i}{2}$ elements of A. We call this a type i iteration. This is possible by Lemma 17.

The resulting set is an ϵ-2-net by construction. It is left to show that $|N| = O_d(\frac{1}{\epsilon^{d-1}})$. In each step of the construction we add at most $\binom{d}{2}$ pairs to N so it is enough to bound the number of iterations T. By the pigeonhole principle, at least $\frac{T}{d-1}$ of the iterations have the same type, say i. There are $\binom{n}{d+1-i}$ subsets of cardinality $d + 1 - i$, and in each of the at least $\frac{T}{d-1}$ type i iterations we 2-stab at least $\Omega_d((\epsilon n)^{d+1-i})$ additional subsets of cardinality $d + 1 - i$, so we have $\frac{T}{d-1} = O_d(\frac{\binom{n}{d+1-i}}{(\epsilon n)^{d+1-i}}) = O_d(\frac{1}{\epsilon^{d+1-i}})$ so $t = O_d(\frac{1}{\epsilon^{d-1}})$ (since $i \geq 2$). This completes the proof of Theorem 16. ◀

Complexity analysis.

The only significant difference between the constructions of Section 4.1 and of Section 4.2 is the factor that depends on the size of the resulting net. Hence, the complexity of the algorithm in this section is bounded by $O_d(n^{2d})$, where d is the 2-VC-dimension of H.

4.2.1   Extension of the Direct Construction of ϵ-2-Nets to ϵ-t-Nets

Now we show how our deterministic construction of ϵ-2-nets can be extended to ϵ-t-nets for other values of t. The following argument is a direct adaptation of the argument from Section 4.2.

▶ **Definition 18.** *Let $H = (V, \mathcal{E})$ be a hypergraph. Given two disjoint vertex sets $A, B \subset V$, we say that A t-stabs B if each hyperedge $e \in \mathcal{E}$ that contains B must contain at least t vertices from A.*

▶ **Proposition 19.** *Let $H = (V, \mathcal{E})$ be a hypergraph with t-VC-dimension d. Then one can construct explicitly an ϵ-t-net for H of size $O_{d,t}(\frac{1}{\epsilon^{d+1-t}})$.*

Proof. Let $S \in \mathcal{E}$ be a hyperedge and let $X \subseteq S$ be a subset of S with cardinality $d + 1$. Since the t-VC-dimension is d then the set X cannot be t-shattered. Notice that $X = X \cap S$ and so each $X' \subset X$ with $|X'| \geq d + 2 - t$ is t-realized by H. For our purpose, we can also assume that $\emptyset$ is t-realized by H (with respect to X), for otherwise the set of all t-subsets in X is a transversal for H of size $\binom{d+1}{t}$. This means that there exists a subset $X \setminus A \subset X$ of size between 1 and $d + 1 - t$ that is not t-realized by H (with respect to X). Equivalently, there is a partition, say $X = A \cup (X \setminus A)$ such that A t-stabs $X \setminus A$. Let $i = |A|$. Note that $i \in \{t, \ldots, d\}$. We say that $X = A \cup (X \setminus A)$ is a type i partition. Note that there could be more than one type partition for the same set X. We need the following lemma:

▶ **Lemma 20.** *Let S be a hyperedge containing $k \geq d+1$ vertices of V. Then there exists an integer $i \in \{t, \ldots, d\}$ and a subset $A \subset S$ with cardinality i that t-stabs $\frac{\binom{k}{d+1}}{(d-t+1)\binom{k}{i}} = \Omega_{d,t}(k^{d+1-i})$ subsets B of cardinality $d+1-i$.*

Proof. For each subset in $\binom{S}{d+1}$ there exists a partition with one of the above stabbing types. By the pigeonhole principle at least $\frac{\binom{k}{d+1}}{d-t+1} = \Omega(k^{d+1})$ of these subsets have the same type, say i. Each such subset X of type i is charged by a t-piercing subset $A \subset X$ of cardinality i. Then by the pigeonhole principle there is a subset $A \subset S$ of cardinality i that is charged at least $\Omega(\frac{k^{d+1}}{\binom{k}{i}}) = \Omega(k^{d+1-i})$ times. This means that A 2-stabs $\Omega(k^{d+1-i})$ subsets $B \subset S$ of cardinality $d+1-i$, as asserted. ◀

Constructing ϵ-t-nets

Let $H = (V, \mathcal{E})$ be as above and let $\epsilon > 0$ be fixed. Put $n = |V|$. We construct an ϵ-t-net of size $O(\frac{1}{\epsilon^{d+1-t}})$ as follows. We start with a set $N = \emptyset$. As long as there is a hyperedge $S \in \mathcal{E}$ with $|S| \geq \epsilon n$ that does not include any t-subset of N, for an appropriate $i \in \{t, \ldots, d\}$ we take an i-subset $A \subset S$ t-stabbing $\Omega((\epsilon n)^{d+1-i})$ subsets of S with cardinality $d+1-i$, and add to N all $\binom{i}{t}$ elements of A. We call this a type i iteration. This is possible by Lemma 20.

Obviously the resulting set is an ϵ-t-net by construction. It is left to show that $|N| = O(\frac{1}{\epsilon^{d+1-t}})$. In each step of the construction we add at most $\binom{d}{t}$ subsets to N so it is enough to bound the number of iterations. Denote this number by L. By the pigeonhole principle, at least $\frac{L}{d-t+1}$ of the iterations have the same type, say i. There are $\binom{n}{d+1-i}$ subsets of cardinality $d+1-i$ and in each of the at least $\frac{L}{d-t+1}$ type i iterations we t-stab at least $\Omega((\epsilon n)^{d+1-i})$ additional subsets of cardinality $d+1-i$. We have that $\frac{L}{d-t+1} = O(\frac{\binom{n}{d+1-i}}{(\epsilon n)^{d+1-i}}) = O(\frac{1}{\epsilon^{d+1-i}})$ so $L = O(\frac{1}{\epsilon^{d+1-t}})$ (since $i \geq t$). This completes the proof. ◀

4.3 t-**VC-dimension versus classical VC-dimension**

What can be said about the relation between VC-dimension and our newly introduced t-VC-dimension, for $t \geq 2$? By definition, $\dim H \leq \dim_2 H$. As shown below ideas from Dudley's unpublished lecture notes [18, Th. 4.37] yield $\dim_2 H \leq 2\dim H + 1$. This is sharp for some small hypergraphs, such as that with vertex set $\{a, b, c\}$ and hyperedges $\{a\}$, $\{b, c\}$, $\{a, c\}$, and $\{a, b, c\}$, which has VC-dimension 1 but 2-VC-dimension 3.

▷ Claim 21. Let $H(V, \mathcal{E})$ be a hypergraph then $\dim_2 H \leq 2\dim H + 1$.

Proof. Assume that V be 2-shattered. We can show that for every $V' \in 2^V$ either V' or $V \setminus V'$ is shattered. This yields the desired result by taking V' of cardinality $\lceil \frac{|V|}{2} \rceil$.

If V' is not shattered, then there exists $S \in 2^{V'} \setminus \Pi_H(V')$. For any $T \in 2^{V \setminus V'}$, there is a set W, $|W| \leq 1$, such that $S \cup T \cup W \in \Pi_H(V)$, because V is 2-shattered. Since $S \notin \Pi_H(V')$ we must have $W \subseteq V'$. But then this implies $T \in \Pi_H(V \setminus V')$, that is, $V \setminus V'$ is shattered. ◀

For general t, we conjecture that $\dim_t H \leq 2\dim H + 2t - 1$. The reasoning below gives roughly $\dim_t H \leq 9.09\max\{\dim H, t-1\}$.

Let H be a hypergraph of finite VC-dimension with a largest t-shattered subset of vertices T. As T is t-shattered, we have $2^T = \{e \setminus S : e \in \Pi_H(T),\ S \subseteq T,\ |S| < t\}$. This yields

$$2^{\dim_t H} \leq |\Pi_H(T)| \cdot \sum_{i=0}^{t-1} \binom{\dim_t H}{i} \leq \sum_{i=0}^{\dim H} \binom{\dim_t H}{i} \cdot \sum_{i=0}^{t-1} \binom{\dim_t H}{i},$$

with the last inequality following from Lemma 4. When $\dim_t H \geq 2\max\{t-1, \dim H\}$, applying Equation (1) gives

$$1 \leq h\left(\frac{\dim H}{\dim_t H}\right) + h\left(\frac{t-1}{\dim_t H}\right).$$

From this inequality we obtain:

▶ **Proposition 22.** *For $t \in \mathbb{N} \setminus \{0\}$, the t-VC-dimension of a hypergraph of VC-dimension d is at least d, at most $2\gamma_2 \max\{d, t-1\}$ (where $\gamma_2 \simeq 4.54$), and, as $d \to \infty$, at most $2d + o(d)$.*

An interesting geometric example is the hypergraph H whose vertex set is a finite subset of $\mathbb{R}^{d-1}$ and whose hyperedges are induced by half-spaces. It is well-known that $\dim H = d$.

More generally, we have $\dim_t H \leq td$ for all t. Indeed, by Tverberg's theorem (see, e.g., [31]), every set T of $td + 1$ points in $\mathbb{R}^{d-1}$ admits a partition into $t + 1$ pairwise disjoint and non-empty sets $T = X \cup Y_1 \cup \cdots \cup Y_t$ such that the intersection of their convex hulls is non-empty. No half-space can t-realize X since any half-space that contains X must contain at least one point from each Y_i, that is, at least t points of $T \setminus X$.

Therefore, for this hypergraph and $t = 2$, the direct construction yields an ϵ-2-net of size $O_d(1/\epsilon^{2d-1})$, while the trivial construction (described at the end of Section 4.1) yields only a weaker upper bound of $O_d(1/\epsilon^{2d})$. With good bounds on $\dim H_{lc}^2$, the construction via H_{lc}^2 (see again Section 4.1) might provide even smaller ϵ-2-nets. In the plane (namely, where $d = 3$), it follows from [21] that $\dim H_{lc}^2 \leq 5$, and so the upper bounds obtained using the direct construction and using H_{lc}^2 are the same – $O(1/\epsilon^5)$.

5   Geometric ϵ-2-Nets

For a fixed $\epsilon > 0$, any hypergraph with VC-dimension d and $n \geq \frac{C_d}{\epsilon^{2^{d+1}}}$ vertices admits, by Theorem 2, an ϵ-2-net of size $O(\frac{d}{\epsilon} \log \frac{1}{\epsilon})$. This leaves open two interesting questions:

1. In cases where the hypergraph admits an ϵ-net of small size, say $O(\frac{1}{\epsilon})$, does it also admit an $O(\frac{1}{\epsilon})$-sized ϵ-2-net (or, more generally, ϵ-t-nets)?
2. Does this extend to smaller values of n?

In this section we answer both in the affirmative for several classes of geometrically-defined hypergraphs.

▶ **Definition 23.** *Given two families B and R of sets, the intersection hypergraph $H(B, R)$ is the hypergraph on vertex set B, where any $r \in R$ defines a hyperedge $\{b \in B : b \cap r \neq \emptyset\}$.*

Note that $H(B, R)$ and $H(R, B)$ are (in general) not isomorphic but dual to each other. Intersection hypergraphs are ubiquitous in discrete and computational geometry. Particular attention is given to the case where either B or R is a set of points, with $H(B, R)$ respectively known as a *primal hypergraph* defined by R or a *dual hypergraph* defined by B. See the survey [33] and the references therein.

We present below several intersection hypergraphs that admit $O(\frac{1}{\epsilon})$-sized ϵ-nets, and prove that each of them has ϵ-2-nets of the same size. Furthermore, while Theorem 2 applies only to hypergraphs with a very large number of vertices, the geometric hypergraphs discussed do not have to contain "many" vertices in order to guarantee the existence of "small" ϵ-2-nets. In some cases (see, e.g., subsection 5.3.1), the behavior is sharp: we can point out two constants $c_1 < c_2$ s.t. if the number of vertices satisfies $|V| \geq \frac{c_2}{\epsilon}$ the hypergraph admits an $O(\frac{1}{\epsilon})$-sized ϵ-2-net, while for $|V| \leq \frac{c_1}{\epsilon}$, there exist hypergraphs from the same family that admit only ϵ-2-nets of size $\Omega(\frac{1}{\epsilon^2})$.

5.1 Non-piercing regions

For our first example we consider a large class of geometric objects introduced by Raman and Ray [37]. A *family of non-piercing regions* is a family of regions of $\mathbb{R}^2$ such that for any two regions γ_1 and γ_2 the difference $\gamma_1 \setminus \gamma_2$ is connected. (Each region may contain holes. See [37] for the exact definitions.)

This extends the more familiar notion of pseudo-disks.

▶ **Theorem 24.** *The intersection hypergraph of two families B and R of non-piercing regions with B finite admits an ϵ-net of size $O(\frac{1}{\epsilon})$ and, if $\epsilon|B| \geq 2$, an ϵ-2-net of size $O(\frac{1}{\epsilon})$.*

The proof relies on several intermediary results. The first one is about an analogue of the Delaunay graph for non-piercing regions [37]. The important specific case where the regions are pseudo-disks had already been studied [5, 25, 26].

▶ **Definition 25.** *A planar support for the hypergraph $(V, \mathcal{E})$ is a planar graph G on the same vertex set V such that any hyperedge in $\mathcal{E}$ induces a connected subgraph of G.*

▶ **Theorem 26** ([37]). *Given two families B and R of non-piercing regions, B finite, their intersection hypergraph $H(B, R)$ admits a planar support.*

The following corollary has already been noted for families of pseudo-discs [5].

▶ **Corollary 27.** *Given two families B and R of non-piercing regions, $\dim H(B, R) \leq 4$.*

Proof. Let $B' \subseteq B$ be a shattered subset of vertices in $H(B, R)$. As the non-piercing property is clearly hereditary, the hypergraph $H(B', R)$ also admits a planar support. For every pair of vertices in B' there exists a hyperedge of $H(B', R)$ that contains these two vertices and no other. Following Definition 25 these two vertices must share an edge in any planar support of $H(B', R)$. Thus said *planar* support is a complete graph on B', forcing $|B'| \leq 4$. ◀

Proof of Theorem 24. First we observe that $H(B, R)$ has ϵ-nets of size $O(\frac{1}{\epsilon})$. Since $H(B, R)$ is finite, we may assume that R is finite as well. To paraphrase from Pyrga and Ray [35, Theorem 4], the following properties suffice:

- For any $0 < \epsilon < 1$ and any $B' \subseteq B$, $H(B', R)$ admits an ϵ-net whose size depends only on ϵ.
- There exist constants $\alpha > 0$, $\beta \geq 0$ and $\tau > 0$ s.t. for any $R' \subseteq R$ there is a graph $G_{R'} = (R', E_{R'})$ with $|E_{R'}| \leq \beta|R'|$ so that for any element $b \in B$ we have $m_b \geq \alpha n_b - \tau$, where n_b is the number of regions of R' intersecting b and m_b is the number of edges in $E_{R'}$ whose both endpoints (which are regions of R') intersect b.

The first condition is verified because $\dim H(B', R) \leq 4$ for every B'. For the second one, let $\alpha = \tau = 1$ and $\beta = 3$, and let $G_{R'}$ be a planar support of $H(R', B)$. (Note the use of duality!) The inequalities follow from its planarity and the connectedness of the subgraph "cut out" by each $b \in B$.

Finally, to obtain an ϵ-2-net, let $K_1 \subseteq B$ be an ϵ-net for $H(B, R)$ of size $O(\frac{1}{\epsilon})$. Let R' consist of the regions of R, if any, that intersect $\geq \epsilon|B|$ regions of B but only one of K_1, and let K_2 be an $\frac{\epsilon}{2}$-net for $H(B \setminus K_1, R')$ also of size $O(\frac{1}{\epsilon})$. Then the desired ϵ-2-net consists of all edges in a planar support of $H(K_1 \cup K_2, R)$. ◀

5.2   Small union complexity

Next, we prove the existence of a small ϵ-2-net for the intersection hypergraph of regions in the plane with linear union complexity and points (i.e. the dual hypergraph defined by the regions).

The union complexity of a family of objects is the function $\kappa : \mathbb{N} \to \mathbb{N}$ that sends each $n \in \mathbb{N}$ to the number of faces of all dimensions in the boundary of the union of $\leq n$ objects, maximized over all subsets of $\leq n$ objects. If $\kappa(n) = O(n)$, we say that the family has *linear union complexity.* Families with linear union complexity include, e.g., families of pseudo-discs: the boundary of the union of $n \geq 3$ pseudo-discs consists of at most $6n - 12$ arcs and as many vertices [23].

The $(\leq k)$-level complexity of the family is defined by counting all faces included in at most k objects (not just on the boundary). To make these definitions precise, one needs to define faces and their dimension; see the survey by Agarwal, Pach and Sharir [3].

A specific case of the following result could also be derived from previous results on Mnets [19], if one adds the additional assumption that the regions have bounded "semi-algebraic description complexity". (The proof of [19] is involved and uses algebraic arguments).

▶ **Theorem 28.** *Let L be a finite family of regions in $\mathbb{R}^2$ with linear union complexity and let $P \subseteq \mathbb{R}^2$ be a set of points. If $|L| \geq \frac{2}{\epsilon}$ then $H(L, P)$ admits an ϵ-2-net of size $O(\frac{1}{\epsilon})$.*

Proof. Let $n := |L|$. First, construct a set $K \subseteq L$ of size $O(\frac{1}{\epsilon})$ such that every "heavy" point of P is included in at least two elements of K, as in the proofs of Theorem 3 or Theorem 24. This relies on the existence of ϵ-nets of size $O(\frac{1}{\epsilon})$ for $H(L, P)$, a result of Aronov, Ezra and Sharir [6].

Since linear union complexity is a hereditary property, K as a subset of L also has linear union complexity. By a standard argument using the Clarkson–Shor theorem [16], the (≤ 2)-level complexity of K is linear as well. Hence, by Euler's formula, the number of hyperedges of size 2 in $H(K, P)$ (whose order of magnitude is equal to the number of (≤ 2)-level faces in the arrangement of K) is at most $c|K|$ for some constant c. By the pigeonhole principle, some region $d \in K$ participates in at most c such hyperedges (i.e., pairs of regions). We pick these at most c pairs of regions to be elements of the ϵ-2-net we construct, and repeat the process for $K \setminus \{d\}$.

We continue in this fashion until all elements of K are removed, and set the ϵ-2-net N to be the set of pairs we picked. Clearly, $|N| = O(|K|) = O(\frac{1}{\epsilon})$. To see that N is indeed an ϵ-2-net, let p be a point that belongs to at least ϵn regions of L. By construction, p belongs to at least two regions of K. Consider the process in which the elements of K are gradually removed, until none of them are left. As a single region is removed at every step, we can look at the step in which the number of remaining regions that contain p is reduced from 2 to 1. Since at that step p is included in exactly two regions of the arrangement, the corresponding pair of regions is added to the ϵ-2-net. Hence, p is covered by both elements of a pair in the ϵ-2-net, as asserted. This completes the proof.   ◀

▶ Remark 29. By essentially the same argument, the hypergraph $H(L, P)$ admits an ϵ-t-net of size $O_t(\frac{1}{\epsilon})$ for any constant $t \leq \epsilon |L|$.

We can extend Theorem 28 to a family L with union complexity $\kappa(n) = n \cdot f(n)$. In this case, the size of the ϵ-2-net is $O(\frac{1}{\epsilon} \cdot \log f(\frac{1}{\epsilon}) \cdot f(\frac{1}{\epsilon} \cdot \log f(\frac{1}{\epsilon})))$. For example, if $\kappa(n) = n \log n$, then one obtains an ϵ-2-net of size $O(\frac{1}{\epsilon} \cdot \log \frac{1}{\epsilon} \cdot \log \log \frac{1}{\epsilon})$.

Indeed, by [6], the hypergraph $H(L, P)$ admits an ϵ-net of size $O(\frac{1}{\epsilon} \cdot \log f(\frac{1}{\epsilon}))$. Let $n' = \frac{1}{\epsilon} \cdot \log f(\frac{1}{\epsilon})$. By the Clarkson–Shor theorem [16], the (≤ 2)-level complexity is bounded
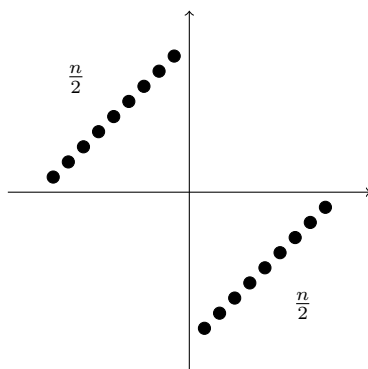
■ **Figure 2** A set P of n points such that in the hypergraph of P with respect to axis-parallel rectangles, for $\epsilon = \frac{2}{n}$, each ϵ-2-net is of size $\Omega(n^2) = \Omega(\frac{1}{\epsilon^2})$.

by $O(n' \cdot f(n'))$, hence there exists a region that participates in at most $f(n')$ hyperedges of order 2. This means that the size of the obtained ϵ-2-net is bounded by $O(n' \cdot f(n'))$.

5.3 More Geometric ϵ-2-Nets

5.3.1 Frames

The next class of intersection hypergraphs we consider is that of points with respect to frames, where a frame is the boundary of an axis-parallel rectangle.

▶ **Proposition 30.** *Let P be a finite set of points of $\mathbb{R}^2$ and let F be a family of frames. If $|P| \geq \frac{5}{\epsilon}$, then $H(P, F)$ admits an ϵ-2-net of size $\leq \frac{8}{\epsilon} - 2$.*

Proof. Let $n = |P|$. If at least ϵn points lie on a same frame, then one of its four sides contains at least $\lceil \frac{\epsilon}{4} n \rceil \geq 2$ of them. Thus it is sufficient to take an $\frac{\epsilon}{4}$-2-net for P with respect to the family of all axis-parallel segments. For any such axis-parallel segment ℓ, take all pairs consisting of the $i\lceil \frac{\epsilon}{4} n \rceil$-th and $(i + 1)\lceil \frac{\epsilon}{4} n \rceil$-th vertices on ℓ. In total, at most $\frac{8}{\epsilon} - 2$ pairs also suffice to pierce all axis-parallel segments. ◀

The interesting behavior here is that the requirement that the vertex set is large cannot be omitted. Consider for example the set P of n points depicted in Figure 2, and let $\epsilon = \frac{2}{n}$. For any pair $\{p_1, p_2\} \subset P$ such that p_1 is in the first quadrant and p_2 is in the third quadrant, there exists a frame r such that $r \cap P = \{p_1, p_2\}$. Hence, any ϵ-2-net for the intersection hypergraph of P and a sufficiently rich family of frames is of size $\geq \frac{n^2}{4} = \frac{1}{\epsilon^2}$.

5.3.2 Axis-parallel rectangles

We conclude this section with the intersection hypergraph $H(P, R)$ of points and axis-parallel rectangles. This hypergraph admits an ϵ-net of size $O(\frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$ [6] and, if $|P| \gg \frac{1}{\epsilon^2}$, an ϵ-2-net of size $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ by Theorem 2. (This last bound follows from the known fact that $\pi^*_{H(P,R)}(m) = \Theta(m^2)$, and from Theorem 2.)

We extend the result to smaller values of $|P|$, at the expense of slightly increasing the size of the obtained ϵ-2-net.

▶ **Theorem 31.** *Let P be a finite set of points in $\mathbb{R}^2$, and let R be a family of axis-parallel rectangles. Assume that $\epsilon|P| \geq 3$. Then $H(P, R)$ admits an ϵ-2-net of size $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$.*

Proof. Let $n := |P|$. First, we construct a set $K \subseteq P$ that intersects every "heavy" rectangle of R at least three times. Let K_1 be an ϵ-net of size $O(\frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$ for $H(P, R)$, whose existence is a result of Aronov, Ezra and Sharir [6]. Let R' consist of those rectangles in R, if any, that contain ϵn points of P but only either one or two points of K_1. They contain at least $\epsilon n - 2 \geq \epsilon n/3$ points of $P \setminus K_1$ which is greater than 1 by our assumption that $\epsilon n \geq 3$. Then let K_2 be an $\epsilon/3$-net for $H(P \setminus K_1, R')$. It hits every rectangle of R' at least once. Finally let R'' consist of the rectangles in R' that contain only two points of $K_1 \cup K_2$ and take an $\epsilon/3$-net K_3 for $H(P \setminus K_1 \cup K_2, R'')$. We let $K = K_1 \cup K_2 \cup K_3$. It contains $k = O(\frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$ points, and each rectangle of R containing ϵn points of P contains at least three points of K. Thus we restrict ourselves to finding a $\frac{3}{k}$-2-net for $H(K, R)$.

As coordinate-wise monotone transformations of the plane do not affect the intersections between K and R, we may assume $K \subset \{0, 1, \ldots, k-1\}^2$ and that the width and height of every rectangle in R are in $\{1, 2, \ldots, k\}$. Define the aspect ratio of a rectangle as the ratio of its height to its width, and for each integer i with $|i| \leq \log k$ let R_i be the family of all axis-parallel rectangles of aspect ratio 2^i. Every such R_i is a family of pseudo-discs: the boundaries of any two of its elements intersect at most twice. Take a $\frac{2}{k}$-2-net of size $O(k)$ for each R_i, altogether $O(k \log k)$ pairs. Each heavy rectangle of R is also the union of two rectangles with the same aspect ratio 2^i for some i, one of which must contain at least two points of K. (This idea is borrowed from Ackerman and Pinchasi [2].)

Hence, the pairs form a $\frac{2}{k}$-2-net of size for each $H(K, R_i)$, a $\frac{3}{k}$-2-net for $H(K, R)$ and an ϵ-2-net for $H(P, R)$. ◀

The case of axis-parallel rectangles illustrates a common phenomenon with ϵ-2-nets: the bounds on the size of ϵ-2-nets worsen as the number of rectangles decreases until, in the smallest case, hypergraphs on only $\frac{2}{\epsilon}$ vertices may require as many as $\Omega(\frac{1}{\epsilon^2})$ pairs.

6   Applications of ϵ-t-Nets

We now describe several contexts in which ϵ-t-nets appear naturally, and present possible applications of our results.

6.1   The Turán problem for hypergraphs

Turán's celebrated theorem in graph theory determines the largest possible number of edges in a graph that does not contain the complete graph K_t (for a fixed integer t) as a subgraph. In 1941, Turán raised a similar question for hypergraphs: the maximum number of hyperedges that a t-uniform hypergraph on n vertices can possess without containing the complete t-uniform hypergraph on k vertices as a sub-hypergraph is the *Turán number* $T(n, k, t)$.

While Turán's theorem for graphs is sharp, the problem for hypergraphs remains notoriously difficult. For $k > t > 2$, there are no known closed expressions for $n \mapsto T(n, k, t)$ (whereas $t = 2$ corresponds to graphs). Determining $T(n, 4, 3)$ is considered one of the major open problems of hypergraph theory [24].

Computing $T(n, k, t)$ amounts to finding the cardinality of a smallest (k/n)-t-net for the complete k-uniform hypergraph on n vertices. Indeed, let $H = (V, \binom{V}{k})$ be said complete hypergraph and let $N \subseteq \binom{V}{t}$. The set N of t-subsets is a (k/n)-t-net for H if and only if every set of k vertices of V contains a t-subset in N. Equivalently, there is no set of k vertices $U \in \binom{V}{k}$ such that $\binom{U}{t} \subseteq \binom{V}{t} \setminus N$. Still equivalently, the complement $\binom{V}{t} \setminus N$ is the set of hyperedges of a t-uniform hypergraph on V that does not contain the complete t-uniform hypergraph on k vertices. We conclude:

▶ **Proposition 32.** *The size of a smallest (k/n)-t-net for the complete k-uniform hypergraph on n vertices is $\binom{n}{t} - T(n, k, t)$.*

6.2 Edge coloring of hypergraphs

Ackerman, Keszegh and Pálvölgyi [1] introduced the problem of coloring t-subsets of vertices in a hypergraph in such a way that each hyperedge contains t-subsets of all colors. They focused on coloring 2-subsets (i.e., edges) in geometric hypergraphs, and in particular on coloring the pairs that are themselves hyperedges of the hypergraph. They obtained constant bounds on the number of colors required for various classes of geometric hypergraphs. One of their main results is the following.

▶ **Theorem 33** ([1, Theorem 4]). *For every dimension d and integers $t \geq 2$, $k \geq 1$ and h, there exists an integer m with the following property: given a set H of h half-spaces in $\mathbb{R}^d$, the t-subsets of every finite set of points S in $\mathbb{R}^d$ can be colored with k colors such that every half-space of H that contains at least m points from S contains a t-subset of points of each of the k colors.*

Using ϵ-2-nets, we obtain a result in the same spirit for all hypergraphs with bounded VC-dimension.

▶ **Proposition 34.** *Let H be a hypergraph on n vertices with VC-dimension d and let $\epsilon \in (\frac{2}{n}, 1)$. Then the pairs of vertices of H can be colored with $\Omega(\frac{n^2 \cdot \epsilon^4}{d}(\log \frac{1}{\epsilon})^{-1})$ colors such that each hyperedge of H of size at least ϵn contains a pair of each color.*

Proof. Consider the corresponding pair hypergraph H^2. To each hyperedge of H of size $\geq \epsilon n$ corresponds a hyperedge of H^2 of size $\Omega(\epsilon^2 n^2)$. Take an $\frac{\epsilon^2}{4}$-net for H^2 consisting of $O(\frac{d}{\epsilon^2} \log \frac{1}{\epsilon})$ pairs of vertices, color all its pairs with one color and remove them, and then repeat the procedure.

We can continue until $\Theta(\epsilon^2 n^2)$ pairs have been colored, which is $\Omega(\frac{n^2 \cdot \epsilon^4}{d}(\log \frac{1}{\epsilon})^{-1})$ steps. All remaining pairs then receive any arbitrary color.   ◀

▶ Remark 35. When there is a hyperedge of size $O(\epsilon n)$, any such coloring has $O(\epsilon^2 n^2)$ colors. This is within a factor $O(\frac{d}{\epsilon^2} \log \frac{1}{\epsilon})$ of our lower bound.

To put this result in a perspective, note that if the pairs of vertices of H are colored in $n^2 \epsilon^4 (\log(1/\epsilon))^{-1}$ colors randomly (i.e., for each pair, we pick a color uniformly, independent of other pairs), then for each hyperedge of size ϵn, the probability that it does not contain a pair of a given color is approximately

$$(1 - n^{-2}\epsilon^{-4}\log(1/\epsilon))^{\epsilon^2 n^2/2} \approx \exp(-\epsilon^{-2}\log(1/\epsilon)/2),$$

which is bounded away from zero. The probability that a hyperedge contains edges of all colors is extremely low, and thus it is not clear that all hyperedges should contain pairs of all colors with positive probability. Hence, the construction of Proposition 34 is stronger than the result obtained by a random coloring.

6.3 χ-Boundedness of graphs

The *chromatic number* $\chi(G)$ of a graph G is the minimum number of colors needed to color the vertices of G such that any two neighboring vertices have different colors. The *clique number* $\omega(G)$ is the size of the largest complete subgraph in G. Obviously, we always have $\chi(G) \geq \omega(G)$. A family $\mathcal{F}$ of graphs is χ-*bounded* if this inequality is "not far from being

tight", namely, if there exists a *binding function* $f\colon \mathbb{N} \to \mathbb{N}$ such that $\chi(G) \le f(\omega(G))$ for every $G \in \mathcal{F}$. On the notion of χ-boundedness, see the recent survey [39].

One of the first results on χ-boundedness is a theorem of Wagon [45] showing that, for any fixed k, the class of graphs that do not have k pairwise disjoint edges as an induced subgraph is χ-bounded. The theorem is proved by induction on k. We present the base case $k = 2$, along with its proof, which is needed for understanding our application.

▶ **Theorem 36.** *If a graph G does not contain two disjoint edges as an induced subgraph, then $\chi(G) \le \frac{\omega(G)(\omega(G)+1)}{2}$.*

Proof. Let A be a complete subgraph of G of size $\omega(G)$. Color each vertex of A with a different color. Consider the vertices of $V \setminus A$ and partition them according to their neighbors in A.

No vertex in $V \setminus A$ is adjacent to all vertices in A, as otherwise it could be added to A to form a larger complete subgraph of G.

For each $a \in A$, let S_a be the set of vertices of $V \setminus A$ adjacent to all vertices of A but a. Note that each S_a is an independent set, as if $(b,c) \in \mathcal{E}$ for some $b, c \in S_a$ then $A \cup \{b,c\} \setminus \{a\}$ is a complete subgraph of G of size larger than A. Hence, the vertices of each S_a can be colored with the color of a, and consequently, setting $S = \bigcup_{a \in A} S_a$, the set $A \cup S$ can be properly colored with $|A|$ colors.

Third, for each pair $\{a,b\} \subseteq A$ let $S_{a,b}$ be the set of vertices of $V \setminus A$ adjacent to neither a nor b. As G does not contain a pair of disjoint edges, each set $S_{a,b}$ is independent, and thus can be colored with a single color. Hence the set $S' = \bigcup_{\{a,b\} \subset A} S_{a,b}$ can be colored in $\binom{|A|}{2}$ colors.

Finally, observe that each $v \in V \setminus A$ belongs either to some S_a or to some $S_{a,b}$. Hence, V can be colored in $|A| + \binom{|A|}{2} = \frac{\omega(G)(\omega(G)+1)}{2}$ colors, as asserted.   ◀

The proof for a general k is an easy inductive argument, which essentially repeats the base step presented above. The binding function f it yields is a polynomial of degree $2(k-1)$.

We observe that the argument used in the proof of Theorem 36 can be improved using an *edge-hitting set* (i.e., a set of edges such that each hyperedge contains at least one of them) for an appropriately chosen hypergraph.

Consider the last step of the proof. The crucial observation it uses is that any $v \in V \setminus (A \cup S)$ belongs to some $S_{a,b}$. However, it seems that many of the vertices belong to many sets $S_{a,b}$ (formally, if v has $|A| - \ell$ neighbors in A, then it belongs to $\binom{\ell}{2}$ sets $S_{a,b}$). Hence, it is plausible that we can "cover" all vertices $v \in V \setminus (A \cup S)$ by a smaller number of sets $S_{a,b}$, and hence, reduce the total chromatic number.

This is achieved by using an edge-hitting set. Let H be the hypergraph whose vertex set is A, and whose hyperedge set is $\{e_v\}_{v \in V \setminus (A \cup S)}$, where $e_v = \{a \in A : (a,v) \notin \mathcal{E}\}$. That is, each $v \in V \setminus (A \cup S)$ induces an hyperedge that consists of all its non-neighbors in A. Let $T = \{(w_i, w_i')\}_{i=1,\ldots,m}$ be an edge-hitting set for H. We claim that in the third step of Wagon's proof, instead of taking all pairs $\{a,b\} \subset A$, it is sufficient to take the m pairs $\{w_i, w_i'\}$. Indeed, let $v \in V \setminus (A \cup S)$. By Wagon's argument, v has at least two non-neighbors in A, and hence, the hyperedge e_v is of size ≥ 2. Hence, it contains some edge $(w_i, w_i') \in T$. By the definition of H, this implies $v \in S_{w_i, w_i'}$. Thus, for each $v \in V \setminus (A \cup S)$ we have $v \in S_{w_i, w_i'}$ for some $(w_i, w_i') \in T$, and so, it is sufficient to color the m sets $S_{w_i, w_i'}$.

This reduces the bound on the chromatic number obtained by Wagon from $\omega + \frac{\omega(\omega-1)}{2}$ to $\omega + m$, where m is the smallest size of an *edge-hitting set* for H.

The notion of ϵ-2-nets can be useful in this context in two ways.

First, ϵ-2-nets can be used to construct a good edge-hitting set, using the strategy of constructing an approximate hitting set from an ϵ-net, pioneered by Brönnimann and Goodrich [12] and followed-up in numerous works (e.g., [17, 20]). One possible way to do this is to consider the hypergraph H_{lc}^2 that corresponds to H (see Section 2), find an $\frac{\epsilon}{2}$-net for it (which is an ϵ-2-nets for H, as was shown in Section 3), and use the algorithm of [12] to leverage them into an approximate hitting set of H_{lc}^2, which is a small edge-hitting set for H. Another possible way is working directly with ϵ-2-nets of H (in cases where this approach is advantageous over working with H_{lc}^2) and transforming them into an approximate edge-hitting set of H, using a variant of the algorithm of [12].

Second, we can use an ϵ-2-net to cover all "sufficiently large" hyperedges of H, which corresponds to coloring all vertices of $V \setminus (A \cup S)$ that have sufficiently many non-neighbors in A. Then, it remains to color vertices that have many neighbors in A, and one may hope that since G does not contain a pair of disjoint edges, they can be colored in a relatively small number of colors.

6.4   Secret sharing

The relevance of ϵ-t-nets to secret sharing was described in the introduction. We would like to add a few remarks:

- The classical objective of secret sharing is that, for some threshold t, the secret may be recovered by any coalition of t members, and none of $t-1$ members. In our variant, we do not require that every large coalition be able to recover the secret, but only a certain upward closed subset of coalitions (e.g., those containing at least $\frac{n}{10}$ members, a representative of each university, and a senior researcher). This natural generalization is known as hierarchical secret sharing [42, 43].
- Departing from classical work on secret sharing, our problem has *two thresholds* – a "necessary threshold" of 6 members, required to obtain the secret, and a "sufficient threshold" which consists of a set of coalitions that must be able to obtain the secret. There is no restriction for the coalitions "between the thresholds"; they may or may not be able to obtain the secret. To the best of our knowledge, such a "two thresholds" scheme has not appeared in the secret sharing literature, although it seems interesting and worthy of study.
- Another deviation from the usual setting of secret sharing is that we seek to minimize the number of divided key shares. This goal is natural in settings where the keys are physical, as in the "scientists problem" described in the introduction, or in secret sharing schemes in which the generation and storage of key shares have a cost.

7   Discussion and Open Problems

A hypergraph H with finite VC-dimension d has ϵ-2-nets of size $O(\frac{d}{\epsilon} \log \frac{1}{\epsilon})$ when n is very large as a function of $\frac{1}{\epsilon}$. This upper bound is the best possible in general, and as we saw in Section 3 may also be best possible even if H admits smaller ϵ-nets. However, we conjecture that in any "reasonable" setting, (including, e.g., all the geometric scenarios discussed in Section 5, and all hypergraphs with hereditarily small ϵ-nets), the existence of an ϵ-net of some order of magnitude, implies the existence of an ϵ-2-net of roughly the same order of magnitude.

Furthermore, we are not aware of any hypergraph in which the dependence of n in $\frac{1}{\epsilon}$ has to be as large as in the assumption of Theorem 2. It may be interesting to extend our results to smaller values of n (as a function of $\frac{1}{\epsilon}$), and to understand whether (as in some of the

geometric cases discussed above), there exists a sharp threshold (as a function of $\frac{1}{\epsilon}$) such that if n is above this threshold, then the hypergraph admits an ϵ-2-net of size $\tilde{O}(\frac{1}{\epsilon})$, but if n is below it, then any ϵ-2-net for the hypergraph contains at least $\Omega(\frac{1}{\epsilon^2})$ pairs.

References

1 Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring Delaunay-edges and their generalizations, 2018. `arXiv:1806.03931`.

2 Eyal Ackerman and Rom Pinchasi. On coloring points with respect to rectangles. *Journal of Combinatorial Theory, Series A*, 120(4):811 – 815, 2013. `doi:10.1016/j.jcta.2013.01.005`.

3 Pankaj K. Agarwal, János Pach, and Micha Sharir. State of the union (of geometric objects). In *Proc. Joint Summer Research Conference on Discrete and Computational Geometry: 20 Years Later, Contemporary Mathematics 452, AMS*, pages 9–48, 2008.

4 Noga Alon, Graham Brightwell, H.A. Kierstead, A.V. Kostochka, and Peter Winkler. Dominating sets in k-majority tournaments. *Journal of Combinatorial Theory, Series B*, 96(3):374 – 387, 2006. `doi:10.1016/j.jctb.2005.09.003`.

5 Boris Aronov, Anirudh Donakonda, Esther Ezra, and Rom Pinchasi. On pseudo-disk hypergraphs, 2018. `arXiv:1802.08799`.

6 Boris Aronov, Esther Ezra, and Micha Sharir. Small-size ϵ-nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing*, 39(7):3248–3282, 2010. `doi:10.1137/090762968`.

7 Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. Polytope approximation and the Mahler volume. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 29–42. SIAM, 2012. `doi:10.1137/1.9781611973099.3`.

8 Patrick Assouad. Densité et dimension. *Annales de l'Institut Fourier*, 33(3):233–282, 1983. `doi:10.5802/aif.938`.

9 Amos Beimel. Secret-sharing schemes: A survey. In Yeow Meng Chee, Zhenbo Guo, San Ling, Fengjing Shao, Yuansheng Tang, Huaxiong Wang, and Chaoping Xing, editors, *Coding and Cryptology - Third International Workshop, IWCC 2011, Qingdao, China, May 30-June 3, 2011. Proceedings*, volume 6639 of *Lecture Notes in Computer Science*, pages 11–46. Springer, 2011. `doi:10.1007/978-3-642-20901-7_2`.

10 Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik–Chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989. `doi:10.1145/76359.76371`.

11 Hervé Brönnimann, Bernard Chazelle, and Jiří Matoušek. Product range spaces, sensitive sampling, and derandomization. *SIAM Journal on Computing*, 28(5):1552–1575, 1999. `doi:10.1137/S0097539796260321`.

12 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 12 1995. `doi:10.1007/BF02570718`.

13 Chris Calabro. *The Exponential Complexity of Satisfiability Problems*. Phd thesis, University of California, San Diego, 2009. URL: `https://escholarship.org/uc/item/0pk5w64k`.

14 Timothy M. Chan. Improved deterministic algorithms for linear programming in low dimensions. *ACM Trans. Algorithms*, 14(3):30:1–30:10, June 2018. `doi:10.1145/3155312`.

15 Bernard Chazelle and Jiří Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *Journal of Algorithms*, 21(3):579–597, 1996. `doi:10.1006/jagm.1996.0060`.

16 Kenneth L. Clarkson and Peter W. Shor. Application of random sampling in computational geometry, II. *Discrete & Computational Geometry*, 4:387–421, 1989. `doi:10.1007/BF02187740`.

17 Kenneth L. Clarkson and Kasturi R. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 01 2007. `doi:10.1007/s00454-006-1273-8`.

18 Richard M. Dudley. Notes on empirical processes. Lecture notes, second printing, 2000.

19   Kunal Dutta, Arijit Ghosh, Bruno Jartoux, and Nabil H. Mustafa. Shallow packings, semialge-
         braic set systems, Macbeath regions, and polynomial partitioning. *Discrete & Computational
         Geometry*, 61(4):756–777, 2019. `doi:10.1007/s00454-019-00075-0`.

20   Guy Even, Dror Rawitz, and Shimon (Moni) Shahar. Hitting sets when the VC-dimension is
         small. *Information Processing Letters*, 95(2):358 – 362, 2005. `doi:10.1016/j.ipl.2005.03.
         010`.

21   Nicolas Grelier, Saeed Gh. Ilchi, Tillmann Miltzow, and Shakhar Smorodinsky. On the
         VC-dimension of convex sets and half-spaces, 2019. `arXiv:1907.01241`.

22   David Haussler and Emo Welzl. Epsilon-nets and simplex range queries. *Discrete & Computa-
         tional Geometry*, 2:127–151, 1987. `doi:10.1007/BF02187876`.

23   Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of Jordan regions
         and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational
         Geometry*, 1:59–71, 1986. `doi:10.1007/BF02187683`.

24   Peter Keevash. *Hypergraph Turán problems*, pages 83–140. Cambridge University Press, 2011.
         `doi:10.1017/CBO9781139004114.004`.

25   Chaya Keller and Shakhar Smorodinsky. Conflict-free coloring of intersection graphs of geomet-
         ric objects. *Discrete & Computational Geometry*, 06 2019. `doi:10.1007/s00454-019-00097-8`.

26   Balázs Keszegh. Coloring intersection hypergraphs of pseudo-disks. *Discrete & Computational
         Geometry*, 10 2019. `doi:10.1007/s00454-019-00142-6`.

27   János Komlós, János Pach, and Gerhard Woeginger. Almost tight bounds for ϵ-nets. *Discrete
         & Computational Geometry*, 7(2):163–173, 02 1992. `doi:10.1007/BF02187833`.

28   Chung L. Liu. *Introduction to Combinatorial Mathematics*. McGraw-Hill, New York, 1968.

29   Jiří Matoušek. Approximations and optimal geometric divide-and-conquer. *Journal of
         Computer and System Sciences*, 50(2):203 – 208, 1995. `doi:10.1006/jcss.1995.1018`.

30   Jiří Matoušek. *Geometric Discrepancy: An Illustrated Guide*. Number 18 in Algorithms and
         Combinatorics. Springer, Berlin, New York, 1999. `doi:10.1007/978-3-642-03942-3`.

31   Jiří Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ,
         USA, 2002. `doi:10.1007/978-1-4613-0039-7`.

32   Nabil H. Mustafa and Saurabh Ray. ϵ-Mnets: Hitting geometric set systems with subsets.
         Discrete & Computational Geometry, 57(3):625–640, 2017. `doi:10.1007/s00454-016-9845-8`.

33   Nabil H. Mustafa and Kasturi Varadarajan. Epsilon-approximations and epsilon-nets. In
         Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and
         Computational Geometry, 3rd Edition*, pages 1241–1267. CRC Press, 2018.

34   János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. *J. Amer. Math.
         Soc.*, 26(3):645–658, 2013. `doi:10.1090/S0894-0347-2012-00759-0`.

35   Evangelia Pyrga and Saurabh Ray. New existence proofs for ϵ-nets. In *Proceedings of the
         Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 199–207, New
         York, NY, USA, 2008. ACM. `doi:10.1145/1377676.1377708`.

36   Yuval Rabani and Amir Shpilka. Explicit construction of a small ϵ-net for linear threshold
         functions. *SIAM Journal on Computing*, 39(8):3501–3520, 2010. `doi:10.1137/090764190`.

37   Rajiv Raman and Saurabh Ray. Planar support for non-piercing regions and applications. In
         Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium
         on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, volume 112 of *LIPIcs*,
         pages 69:1–69:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. URL: `http:
         //www.dagstuhl.de/dagpub/978-3-95977-081-1`, `doi:10.4230/LIPIcs.ESA.2018.69`.

38   Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*,
         13(1):145 – 147, 1972. `doi:10.1016/0097-3165(72)90019-2`.

39   Alex Scott and Paul Seymour. A survey on χ-boundedness, 2018. `arXiv:1812.07500`.

40   Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
         `doi:10.1145/359168.359176`.

41   Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J. Math.*, 41(1):247–261, 1972. URL: `https://projecteuclid.org:443/euclid.pjm/1102968432`.

42   Gustavus J. Simmons. How to (really) share a secret. In Shafi Goldwasser, editor, *Advances in Cryptology - CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings*, volume 403 of *Lecture Notes in Computer Science*, pages 390–448. Springer, 1988. `doi:10.1007/0-387-34799-2_30`.

43   Tamir Tassa. Hierarchical threshold secret sharing. *Journal of Cryptology*, 20(2):237–264, 2007. `doi:10.1007/s00145-006-0334-8`.

44   Vladimir N. Vapnik and Alexei Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

45   Stanley Wagon. A bound on the chromatic number of graphs without certain induced subgraphs. *Journal of Combinatorial Theory, Series B*, 29(3):345 – 346, 1980. `doi:10.1016/0095-8956(80)90093-3`.

46   Emo Welzl. Partition trees for triangle counting and other range searching problems. In *Proceedings of the Fourth Annual Symposium on Computational Geometry (Urbana, IL, 1988)*, pages 23–33. ACM, New York, 1988. `doi:10.1145/73393.73397`.