


On the VC-dimension of convex sets and half-spaces

Nicolas Grelier ^{*} Saeed Gh. Ilchi ^{*} Tillmann Miltzow [†]
 Shakhar Smorodinsky[‡]

Abstract

A family S of convex sets in the plane defines a hypergraph $H = (S, \mathcal{E})$ as follows. Every subfamily $S' \subset S$ defines a hyperedge of H if and only if there exists a halfspace h that fully contains S' , and no other set of S is fully contained in h . In this case, we say that h realizes S' . We say a set S is shattered, if all its subsets are realized. The VC-dimension of a hypergraph H is the size of the largest shattered set.

We show that the VC-dimension for *pairwise disjoint* convex sets in the plane is bounded by 3, and this is tight. In contrast, we show the VC-dimension of convex sets in the plane (not necessarily disjoint) is unbounded. We also show that the VC-dimension is unbounded for pairwise disjoint convex sets in \mathbb{R}^d , for $d \geq 3$. We focus on, possibly intersecting, segments in the plane and determine that the VC-dimension is always at most 5. And this is tight, as we construct a set of five segments that can be shattered. We give two exemplary applications. One for a geometric set cover problem and one for a range-query data structure problem, to motivate our findings.



The indicated halfspace realizes the hyperedge $\{2, 4\}$.

1 Introduction

Geometric hypergraphs (also called range-spaces) are central objects in computational geometry, statistical learning theory, combinatorial optimization, linear programming, discrepancy theory, data bases and several other areas in mathematics and computer science.

In most of these cases, we have a finite set P of points in \mathbb{R}^d and a family of simple geometric regions, such as say, the family of all halfspaces in \mathbb{R}^d . Then we consider the combinatorial structure of the set system $(P, \{h \cap P\})$ where h is any halfspace. Many optimization problems can be formulated on such structures. A key property that such hypergraphs have is the so-called bounded VC-dimension (see below for exact definitions). In this paper we initiate the study of a more complicated structure by allowing the underlying set of vertices to be arbitrary convex sets and not just points. We show that when the underlying family consists of pairwise disjoint convex sets in the plane then the corresponding hypergraph

^{*}Department of Computer Science, ETH Zürich, Switzerland

[†]Department of Computer Science, Utrecht University, work supported by Veni grant EAGER by the NWO.

[‡]Department of Mathematics, Ben-Gurion University of the NEGEV, Beer-Sheva, Israel . Work supported by Grant 635/16 from the Israel Science Foundation.

has VC-dimension at most 3 and this is tight. We also show that when the sets may have intersection then the VC-dimension is unbounded. We also show that even for pairwise disjoint convex sets in \mathbb{R}^d the VC-dimension is unbounded already for $d \geq 3$.

We note that many deep results that hold for arbitrary hypergraphs with bounded VC-dimension readily apply to such hypergraphs. This includes, e.g., bounds on the discrepancy of such hypergraphs, bounds of $O(\frac{1}{\varepsilon^2})$ on the size of ε -approximations and also bounds on matchings or spanning trees with (so-called) low crossing numbers (see, e.g., [4, 7, 9, 14]).

Preliminaries and Previous Work A hypergraph $H = (V, \mathcal{E})$ is a pair of sets such that $\mathcal{E} \subseteq 2^V$. A geometric hypergraph is one that can be realized in a geometric way. For example, consider the hypergraph $H = (V, \mathcal{E})$, where V is a finite subset of \mathbb{R}^d and \mathcal{E} consists of all subsets of V that can be cut-off from V by intersecting it with a shape belonging to some family of “nice” geometric shapes, such as the family of all halfspaces. See Figure 1, for an illustration of a hypergraph induced by points in the plane with respect to discs.




Figure 1: This is a representation of the hypergraph $H = (V, \mathcal{E})$, with $V = \{1, 2, 3, 4\}$ and \mathcal{E} contains all subsets of V , except $\{2, 4\}$. The disks realizing all possible hyperedges of size 2 are indicated.

The elements of V are called *vertices*, and the elements of \mathcal{E} are called *hyperedges*.

We consider the following kinds of geometric hypergraphs: Let \mathcal{C} be a family of convex sets in \mathbb{R}^2 (or, in general, in \mathbb{R}^d). We say that a subfamily $S \subseteq \mathcal{C}$ is *realized* if there exists a halfspace h such that $S = \{C \in \mathcal{C} \mid C \subset h\}$. In words, there exists a halfspace h such that the subfamily of \mathcal{C} of all sets that are fully contained in h is exactly S . We refer to the hypergraph $H = (\mathcal{C}, \{S \mid S \text{ is realized}\})$ as the hypergraph induced by \mathcal{C} . In the literature, hypergraphs that are induced by points with respect to geometric regions of some specific kind are also referred to as *range spaces*. We start by introducing the concept of VC-dimension.

VC-dimension and ε -nets A subset $T \subset V$ is called a *transversal* (or a *hitting set*) of a hypergraph $H = (V, \mathcal{E})$, if it intersects all sets of \mathcal{E} . The *transversal number* of H , denoted by $\tau(H)$, is the smallest possible cardinality of a transversal of H . The fundamental notion of a transversal of a hypergraph is central in many areas of combinatorics and its relatives. In computational geometry, there is a particular interest in transversals, since many geometric problems can be rephrased as questions on the transversal number of certain hypergraphs. An important special case arises when we are interested in finding a small size set $N \subset V$ that intersects all “relatively large” sets of \mathcal{E} . This is captured in the notion of an ε -net for a hypergraph:

Definition 1 (ε -net). *Let $H = (V, \mathcal{E})$ be a hypergraph with V finite. Let $\varepsilon \in [0, 1]$ be a real number. A set $N \subset V$ (not necessarily in \mathcal{E}) is called an ε -net for H if for every hyperedge $S \in \mathcal{E}$ with $|S| \geq \varepsilon|V|$ we have $S \cap N \neq \emptyset$.*

In other words, a set N is an ε -net for a hypergraph $H = (V, \mathcal{E})$ if it stabs all “large” hyperedges (i.e., those of cardinality at least $\varepsilon|V|$). The well-known result of Haussler and Welzl [5] provides a combinatorial condition on hypergraphs that guarantees the existence of small ε -nets (see below). This requires the following well-studied notion of the Vapnik-Chervonenkis dimension [13]:

Definition 2 (VC-dimension). *Let $H = (V, \mathcal{E})$ be a hypergraph. A subset $X \subset V$ (not necessarily in \mathcal{E}) is said to be shattered by H if $\{X \cap S : S \in \mathcal{E}\} = 2^X$. The Vapnik-Chervonenkis dimension, also denoted the VC-dimension of H , is the maximum size of a subset of V shattered by H .*

Relation between ε -nets and the VC-dimension Haussler and Welzl [5] proved the following fundamental theorem regarding the existence of small ε -nets for hypergraphs with small VC-dimension.

Theorem 3 (ε -net theorem). *Let $H = (V, \mathcal{E})$ be a hypergraph with VC-dimension d . For every $\varepsilon \in (0, 1)$, there exists an ε -net $N \subset V$ with cardinality at most $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$.*

In fact, it can be shown that a random sample of vertices of size $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ is an ε -net for H with a positive constant probability.

Many hypergraphs studied in computational geometry and learning theory have a “small” VC-dimension, where by “small” we mean a constant independent of the number of vertices of the underlying hypergraph. In general, range spaces involving semi-algebraic sets of constant description complexity, i.e., sets defined as a Boolean combination of a constant number of polynomial equations and inequalities of constant maximum degree, have finite VC-dimension. Halfspaces, balls, boxes, etc. are examples of ranges of this kind; see, e.g., [8, 10] for more details.

Thus, by Theorem 3, these hypergraphs admit “small” size ε -nets. Kórmlos *et al.* [6] proved that the bound $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ on the size of an ε -net for hypergraphs with VC-dimension d is best possible. Namely, for a constant d , they construct a hypergraph H with VC-dimension d such that any ε -net for H must have size of at least $\Omega\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$. Recently, several breakthrough results provided better (lower and upper) bounds on the size of ε -nets in several special cases [1, 2, 11].

In summary, the VC-dimension is a central notion in many areas. It proofed a useful concepts with many applications. To the best of our knowledge the VC-dimension has not been studied for the geometric hypergraphs introduced in this paper.

Results We look at a selection of natural geometric hypergraphs that arise in our setting. Our main contribution is to determine its VC-dimension precisely, in all cases that we consider.

Theorem 4. *Consider a family of sets and the ranges defined by halfspaces as described above.*

1. *For convex sets in the plane, possibly intersecting, the VC-dimension is unbounded.*
2. *For convex disjoint sets in \mathbb{R}^d , for $d \geq 3$, the VC-dimension is unbounded.*
3. *For convex disjoint sets in \mathbb{R}^2 the VC-dimension is at most three and this is tight.*
4. *For segments in the plane, possibly intersecting, the VC-dimension is five and this is tight.*

If we allow our convex sets to intersect, we can construct set systems with unbounded VC-dimension. This shows us that, we need to consider more restricted settings to have bounded VC-dimension. One direction is to restrict ourselves to disjoint convex sets. This leads to VC-dimension three in the plane, but gives still unbounded VC-dimension in higher dimensions. Another approach is to restrict ourselves to simpler geometric objects. Here we focus on line segments. We can show that the VC-dimension for line segments is precisely five.

In summary, we showed that the VC-dimension is unbounded in our setting. We considered two natural restrictions: (a) disjointness of the family of sets. (b) simpler geometric sets. We fully settle the VC-dimension for all of those cases.

On the way to prove Theorem 4, we gain some structural insight, which we believe to be useful also for future work on the topic. First, Assume that (S, \mathcal{E}) is shatterable. We can show that every set $s \in S$ must have at least one vertex on the boundary of the convex hull of $\bigcup_{s \in S} s$, see Lemma 10. Second, for most

hyperedges $S' \subseteq S$, we can identify a canonical halfplane that realizes S' . This gives us a combinatorial way to bound the number of realizable hyperedges, see Lemma 12.

In order to show the relevance of our findings to the field of algorithms, we give two simple exemplary applications that follow easily together with previous work.

Algorithmic Applications To demonstrate the usefulness of VC-dimension, we describe here two algorithmic applications that follow from our results and the machinery developed in the last decades.

For our first expository application, we consider a natural hitting set problem.

Definition 5 (Hitting Halfplanes with segments). *Given a set H of halfplanes and a set S of segments, the halfspace-segment-hitting set problem asks for a minimum set $T \subset S$, such that every halfplane $h \in H$ contains at least one segment $t \in T$ entirely. This is an optimization problem, where we try to minimize the size of T .*

Using the framework of Brönniman and Goodrich [3], we get the following theorem.

Corollary 6. *There is an $O(\log c)$ -approximation algorithm for the halfspace-segment-hitting set problem, where c is the size of the optimal solution.*

Proof. We use Theorem 4 for segments in the plane and the framework from [3]. □

As a second expository application, we have to introduce the problem of approximate range counting. Given a family of sets O and a halfplane h , we denote by $n(h, O)$ the number of sets fully contained in h . We denote by $r(h, O) = \frac{n(h, O)}{|O|}$ the relative number of sets. We want to construct a datastructure that reports a number t such that

$$|r(h, O) - t| < \varepsilon,$$

for some given ε . Thus, here we allow an absolute error rather than a relative error. A simple way is to construct a datastructure is to sample a family of sets $P \subseteq O$ and query how many objects of P are fully contained inside h . If this set P is small we can do queries fast.

Corollary 7. *Let O be a set of disjoint convex objects in the plane. Then there exists a set $P \subseteq O$ of size $O(\frac{1}{\varepsilon^2})$ such that $|r(h, O) - r(h, P)| < \varepsilon$.*

Proof. We use Theorem 4 for disjoint convex sets in the plane and the results from [14]. □

Note that there are many different notions of approximate range counting and we presented here a simple one. Recall that we only want to highlight the relevance of our findings for algorithmic applications.

Structure In Section 2, we show Item 1 and Item 2 of Theorem 4. In Section 3, we handle the case of disjoint sets in the plane, which shows Item 3 of Theorem 4. In Section 4, we show Item 4 of Theorem 4. In Section 5, we will consider the minimum number of intersections that shattered families of sets in the plane must have.

2 Convex sets in the plane and higher dimensions

In this section, we show that when the underlying convex sets may intersect, the VC-dimension can be unbounded.

Lemma 8. *For any $n > 0$, there exists n convex sets in the plane that can be shattered. Furthermore, there exists n pairwise disjoint convex sets in \mathbb{R}^d that can be shattered, for any $d \geq 3$.*

Proof. We denote $[n] = \{1, 2, \dots, n\}$. We place $2^n - 2$ points on the unit circle in the plane as follows. For every non-trivial subset $I \subset [n], I \neq \emptyset, I \neq [n]$ we place a point p_I on that unit circle. For each $j \in [n]$ we define the convex set C_j as the convex hull of all points p_I for which $j \in I$. Namely $C_j = \text{conv}(\{p_I \mid j \in I\})$. We claim that the family $\mathcal{C} = \{C_1, \dots, C_n\}$ is shattered. To see this, let $S \subseteq \mathcal{C}$. If S is either empty or the whole family \mathcal{C} then it is easy to see that it is realized as there is a halfplane containing all sets and there is also a halfplane containing none of the sets. So let I be the corresponding non-trivial set of indices corresponding to the members of S . Put $J = [n] \setminus I$ take a line ℓ separating the point p_J from all




Figure 2: Illustration of the construction, of unbounded VC-dimension, for the case $n = 4$. The halfspace to the right of ℓ realizes the hyperedge $\{C_3\}$.

other points, see Figure 2 for an illustration. We claim that the halfplane h bounded by ℓ and containing those points realizes the subfamily S . Indeed notice that for each $C_i \in S$ all points p_K for which $i \in K$ are contained in h so their convex hull C_i is also contained in h . Note also that for any $C_j \notin S$ we have that $j \in J$ so C_j contains the point p_J and hence it is not fully contained in h . This shows that S is realized for any S and hence \mathcal{C} is shattered.

The VC-dimension is also unbounded for $d \geq 3$ even for pairwise disjoint convex sets. Map each point of C_i such as (x, y) to a point (x, y, i) from \mathbb{R}^2 to \mathbb{R}^3 . With this mapping, all the convex sets will be disjoint and we can still shatter these sets as before, by considering vertical halfspaces. See Figure 3 for an illustration. The case for $d > 3$ follows in the same way. \square

3 Disjoint convex sets in the plane

Here we prove the following lemma.

Lemma 9. *Let \mathcal{C} be a family of pairwise disjoint convex sets in the plane \mathbb{R}^2 . Then, the hypergraph induced by \mathcal{C} has VC-dimension at most 3.*

Note that it is easy to find *three* disjoint convex sets that can be shattered. We first prove the following useful lemma.

Lemma 10 (Convex Hull). *Let \mathcal{C} be a family of sets in the plane. If \mathcal{C} is shattered, then each set in \mathcal{C} contains a point on the boundary of the convex hull of \mathcal{C} .*

Note that in Lemma 10 the sets need not to be convex.

Proof. Let us assume to the contrary that there exists a set C contained in the convex hull of \mathcal{C}' , where \mathcal{C}' is a proper subset of \mathcal{C} . Then any halfplane containing all elements in \mathcal{C}' must also contain C . Therefore, it is not possible to realize \mathcal{C}' , which implies that \mathcal{C} is not shattered. \square

Proof of Lemma 9. Let us assume by contradiction that there exists a shattered family $\mathcal{C} = \{1, 2, 3, 4\}$ of four disjoint convex sets. For each convex set i in \mathcal{C} , we denote by p_i a point in i that lies on the boundary of the convex hull of \mathcal{C} . The existence of this point is assured by Lemma 10. Without loss of generality, let us assume that p_2 and p_4 have the same y -coordinate, p_2 to the left of p_4 , with p_1 above




Figure 3: The sets of the two dimensional construction stacked disjointly on top of one another.




Figure 4: Notation used in the proof of Lemma 9.

and p_3 below them. By assumption, there exist a halfplane h containing 2 and 4 but not 1 nor 3. In particular, h contains p_2 and p_4 . However, as p_1 and p_3 are on the boundary of the convex hull of \mathcal{C} , h must contain at least one of them, say p_1 . We denote by ℓ the bounding line of h , which is therefore below the segment between the points p_2 and p_4 . As the set $\{2, 4\}$ is realized by h , 1 must contain a point q below ℓ . We denote by s the segment between the points p_1 and q . Likewise, we denote by t the segment between the points q and p_3 . Finally, we denote by r the union of s and t . As 1 is convex, s is fully contained inside 1, as its endpoints are contained in 1. By assumption, all i are pairwise disjoint. Thus 2 is not intersecting with s , therefore all points in 2 are to the left of r , as t lies below ℓ and 2 is above ℓ . By the same argument, all points in 4 are to the right of r . However, any halfplane realizing $\{1, 3\}$ must contain r , and would therefore also contain 2 or 4, which is a contradiction. \square

4 Segments

A line segment in the plane can be viewed as the simplest convex set that is not a point. We now turn to study the special case of the VC-dimension of hypergraphs induced by line segments.

Lemma 11. *Let S be a set of (not necessarily disjoint) line segments in \mathbb{R}^2 . Then the hypergraph induced by S has VC-dimension at most 5.*




Figure 5: The halfspace h realizing some set S' indicated in turquoise. The line ℓ bounds h . First, we move ℓ up, and then we rotate it around p_2 and p_3 in that order respectively.

Before proceeding with the proof we need the following lemma. We say that a set of segments is *in general position*, if no three endpoints are collinear.

Lemma 12. *Let S be a set of n segments in the plane, in general position. Then the number of non-empty subsets of S that are realized is at most $2n(n-1) + 2$.*

Proof. Let h be a halfspace realizing a subset $S' \subseteq S$, with $S' \neq \emptyset$ and $S' \neq S$. In the first step, we identify a unique tangent line ℓ , by some transformation argument. In the second step, we show that every pair of segments has at most four tangent lines. Thus, together with the trivial subsets of S , we can realize at most

$$4 \binom{n}{2} + 2 = 2n(n-1) + 2$$

subsets S' .

For the first step, let $S' \subseteq S$, with $S' \neq \emptyset$ and $S' \neq S$. Furthermore, let h be a halfspace that realizes S' . See Figure 5 for an illustration. We denote by ℓ the bounding line. We orient ℓ such that S' lies to the left of ℓ . In this way, ℓ has a head and a tail. If there are several points on ℓ then we can clearly say, which is closest to its head in the obvious way. Translate h inward until its boundary line ℓ hits one element of S' . This must happen as $S' \neq \emptyset$. As the set S is in general position, ℓ touches S' in at




Figure 6: Given two segments there are at most four configurations that can arise. All of them are displayed. Note that if the two segments are not crossing then only two configurations are possible.

most two endpoints p, q . Suppose that, we touch indeed two points, the other case is handled similarly. Furthermore, we say that q is the point closer to the head of ℓ . Then we rotate ℓ counterclockwise around q , up until one of two events happen.

- (a) The line ℓ touches another vertex of some segment $s \in S'$ at its head.
- (b) The line ℓ touches an endpoint of some segment $s \in (S \setminus S')$ at its tail.

Note that it could also be that ℓ touches a vertex of some segment $s \in (S \setminus S')$ at its head. We ignore that case, as this event does not change whether h realizes S' or not. It is easy to see that it is impossible that ℓ touches another vertex of some segment $s \in S'$ at its tail. In case (a), we touch a new point q' and we proceed as before. In other words, we rotate ℓ counterclockwise around q' , up until, either (a) or (b) will happen. In case (b), we stop. Note that since $S' \neq S$, this will eventually happen. We will end up in a configuration, where ℓ touches a vertex of a segment $s \in S'$ at its head and a vertex of another segment $t \in (S \setminus S')$ at its tail. Note that both segments are to the left of ℓ , with respect to the orientation of ℓ . Note that the halfspace h defined by ℓ only needs an infinitesimally small rotation to realize the original set S' that we started with. Thus if there were any halfspace realizing S' , there must be one of the special type, that we just described. This shows the first step. In the second step, we will upper bound the number of those special configurations.

For the second step, consider two segments $a, b \in S$. See Figure 6 for an illustration. Note first that they are either crossing or they are disjoint. One of them must be contained in the set S' that we want to realize and the other is not. This also immediately tells us the orientation of the line in the configuration. It is easy to check that all four configurations are displayed in Figure 6. \square

Proof of Lemma 11: Let S be a set of n line segments that can be shattered. We can assume that S is in general position, by some standard perturbation arguments. We will use the fact that the number of distinct subsets of S that are realized is at most $2n(n-1)+2$, see Lemma 12. As there are 2^n subsets that need to be realized, for S to be shattered, we can conclude that It follows that $2^n \leq 2n(n-1)+2$. However, this inequality is violated for $n \geq 6$, so $n \leq 5$. \square

The next lemma shows the second part of Lemma 11

Lemma 13. *There exists a set of five segments that are shattered by halfplanes.*




Figure 7: Shattering five segments as in Lemma 13. The shattered set is marked turquoise, the other segments are coloured orange. Due to the symmetry of the set of segments, we have shown that all hyperedges can be realized and thus the set can be shattered.

Proof. The set is shown in Figure 7. The segments that are realized are shown in turquoise, and the other in orange. It is easy to find a halfplane realizing none or all the segments. We show in Figure 7 how to realize all the remaining configurations. \square

5 Number of intersections

From Lemma 9 we have that any shattered set of n convex sets are not pairwise disjoint when $n \geq 4$. We show that there are quadratically many pairs of intersecting convex sets.

Lemma 14. *In a shattered set of n convex sets there are at least $n(n-3)/6$ intersections.*

Proof. Consider the intersection graph G of the convex sets. As for any four vertices there is an edge, we obtain that the independence number of G is at most 3. (The independence number of a graph denotes the size of the largest independent set of the graph.) Therefore there is no K_4 in the complement of G . Turán's theorem states that any graph with n vertices not containing K_k has at most $(1 - 1/k) \cdot n^2/2$ edges [12]. Therefore, there are at most $(1 - 1/3) \cdot n^2/2 = n^2/3$ non-edges in G . This is equivalent to having at least $\binom{n}{2} - n^2/3 = n(n-3)/6$ edges in G . \square

It would be interesting to find an upper bound on how few intersections there may be in a shattered set of n convex sets. The question can also be asked for $n \leq 5$ when considering the more specific case of segments. We have given in Lemma 13 a shattered set of five segments with five intersections. We produce now an example of a shattered set with four segments having only one intersection.




Figure 8: Shattering a set of 4 segments, with only one intersection. We indicate how to shatter the sets $\{1, 2\}$, $\{1, 3\}$, $\{1, 2, 3\}$ and $\{1\}$. All other sets follow the same principle.

Lemma 15. *There exists a shattered set of four segments with only one intersection.*

Proof. We consider the four segments as in Figure 8, denoted by $\{1, 2, 3, 4\}$. It is easy to realize none or all segments. To realize three of them, say $\{1, 2, 3\}$ take a halfplane with vertical bounding line immediately after the left end of 4. Likewise to realize two consecutive segments, say $\{1, 2\}$, take a halfplane with diagonal bounding line just above the left end of 3 and to the right of the left end of 4. For opposite segments, say $\{1, 3\}$, take a halfplane with vertical bounding line such that this line intersects both 2 and 4. Finally the reader can check that for each set with exactly one segment i , it is possible to find a halfplane containing only i . \square

Acknowledgement

This work was initiated during the 17th Gremo Workshop on Open Problems 2019. The authors would like to thank the other participants for interesting discussions during the workshop. We are also grateful to the organizers for the invitation and a productive and pleasant working atmosphere.

References

- [1] N. Alon. A non-linear lower bound for planar epsilon-nets. *Discrete & Computational Geometry*, 47(2):235–244, 2012. URL: <http://dx.doi.org/10.1007/s00454-010-9323-7>, doi:10.1007/s00454-010-9323-7.
- [2] B. Aronov, E. Ezra, and M. Sharir. Small-size epsilon-nets for axis-parallel rectangles and boxes. *SIAM J. Comput.*, 39(7):3248–3282, 2010.
- [3] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.
- [4] B. Chazelle and E. Welzl. Quasi-optimal range searching in space of finite vc-dimension. *Discrete & Computational Geometry*, 4:467–489, 1989.
- [5] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.
- [6] J. Komlós, J. Pach, and G.J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.
- [7] J. Matousek. Tight upper bounds for the discrepancy of half-spaces. *Discrete & Computational Geometry*, 13:593–601, 1995.
- [8] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.
- [9] J. Matousek, E. Welzl, and L. Wernisch. Discrepancy and approximations for bounded vc-dimension. *Combinatorica*, 13(4):455–466, 1993.
- [10] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. Wiley Interscience, New York, 1995.
- [11] J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. In *Symposium on Computational Geometry*, pages 458–463, 2011.
- [12] P. Turán. On an external problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941.
- [13] V. N. Vapnik and A. 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.
- [14] Y. Li, P.M. Long, and A. Srinivasan. Improved bounds on the sample complexity of learning. *J. Comput. Syst. Sci.*, 62(3):516–527, 2001. URL: <http://dx.doi.org/10.1006/jcss.2000.1741>, doi:10.1006/jcss.2000.1741.