

UC Davis

UC Davis Electronic Theses and Dissertations

Title

Halving point configurations; techniques from algebraic and convex geometry

Permalink

<https://escholarship.org/uc/item/2c89s4ts>

Author

Leroux, Brett Elliott

Publication Date

2021

Peer reviewed|Thesis/dissertation

Halving point configurations; techniques from algebraic and convex geometry

By

BRETT ELLIOTT LEROUX
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

MATHEMATICS

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Luis Rademacher, Chair

Eric Babson

Jesús De Loera

Committee in Charge

2021

© Brett Elliott Leroux, 2021. All rights reserved.

Contents

Abstract	iv
Acknowledgments	vi
Chapter 1. Introduction	1
Overview of Chapter 2: Preliminaries	3
Overview of Chapter 3: The k -set problem for algebraic set systems	3
Overview of Chapter 4: Generally k -neighborly embedded manifolds and varieties	4
Overview of Chapter 5: Improved bounds for the expected number of k -sets	5
Overview of Chapter 6: Translations of a fixed convex body in the plane	6
Chapter 2. Preliminaries	7
2.1. Discrete geometry and convex geometry	7
2.2. k -sets and k -facets	8
2.3. Mass partitions	9
2.4. Set systems and VC-theory	12
2.5. Classical real algebraic geometry	13
Chapter 3. The k -set problem for algebraic set systems	16
3.1. Introduction	16
3.2. Preliminaries	18
3.3. Counting k -facets via maps	21
Chapter 4. Generally k -neighborly embedded manifolds and varieties	29
4.1. Introduction	29
4.2. Generally neighborly embeddings	30
4.3. Generally neighborly manifolds	32

4.4.	Generally k -neighborly algebraic varieties	33
4.5.	Weakly k -neighborly sets	34
4.6.	Weakly k -neighborly varieties	36
4.7.	Additional evidence	41
4.8.	Neighborly embeddings	42
Chapter 5.	Improved bounds for the expected number of k -sets	44
5.1.	Expected number of k -edges	44
5.2.	Bounding the expected number of k -edges	47
5.3.	On the number of k -edges via the polynomial method	51
Chapter 6.	Translations of a fixed convex body in the plane	59
6.1.	Introduction	60
6.2.	Upper bound for T_C - k -sets, probabilistic, k proportional to n	61
6.3.	Lower bound for T_C - k -sets, deterministic, k proportional to n	64
6.4.	Bounds on the growth function	66
6.5.	Lower bound for T_C - k -sets, probabilistic, some k proportional to n	68
Chapter 7.	Open questions	72
7.1.	The <i>real degree</i> of algebraic varieties	72
7.2.	Conjecture on generally k -neighborly embeddings/manifolds	73
7.3.	Lower bounds for the algebraic k -set problem	73
7.4.	Bound on the number of points of intersection of $Z(f)$ and the k -edge graph	73
7.5.	Polynomial partitioning for k -sets in higher dimensions	74
	Bibliography	75

Abstract

A *halving line* of a set of points is a line that divides the set of points into two equal parts. The *halving lines problem* asks: What is the maximum number of distinct halving lines that a set of n points can have? The focus of this dissertation is on results either about or inspired by the halving lines problem and its variations and generalizations. We start out by generalizing the halving lines problem in the most natural way: Given a family of curves or surfaces and a set of points, we want to know how many ways there are to divide the set of points into two equal parts using one of the curves or surfaces in the given family. And we would also like to know what the maximum number of halving curves or surfaces that a set of n points can have is. This type of problem leads us to ask several new questions which are relevant to discrete geometry, convex geometry, as well as real algebraic geometry. Some of our main contributions are as follows:

- We study a variation on the halving lines problem when the family of separating curves or surfaces is a parametric family of algebraic curves or surfaces. In some cases, we are able to exactly count the number of halving curves. An example when we obtain an exact count is for the conic sections. These results are similar to a result of Ardila on halving circles.
- The concept of neighborliness is crucial for several of our results. Neighborly polytopes are important to the theory of convex polytopes because of their appearance in the upper bound theorem of McMullen. The *moment curve* is the standard way to construct neighborly polytopes. We define *generally neighborly* manifolds and algebraic varieties. These objects can be seen as higher-dimensional analogues of the moment curve.
- We study the random version of the original k -set problem in the plane and establish an improved upper bound for the expected number of k -sets. We also investigate how it may be possible to improve our bound using the continuous version of the polynomial partitioning theorem of Guth and Katz. This motivates a question about the points of intersection of an algebraic curve and the k -edge graph of a set of points.
- Another variation on the random version of the k -set problem is introduced and essentially solved: We obtain nearly tight bounds for the expected number of ways one can enclose k points from a random set of points using a translation of a fixed strictly convex body in

the plane. The motivation is to show that a technique for counting k -sets due to Bárány and Steiger is nearly tight for a natural variation on the k -set problem.

A theme throughout this work is the investigation of questions whose answers help us understand the limits of an argument or proof technique. Most of the ideas presented here also appeared in papers coauthored with Luis Rademacher.

Acknowledgments

First, I would like to thank Luis Rademacher for working with me for the past 3.5 years. I am grateful for his ability to foster math research spanning disparate areas. And I am grateful for his willingness to entertain and stimulate ideas, even those which appear far-fetched.

Many thanks to Jesús De Loera who has been a source of advice and inspiration throughout my time in graduate school. I would also like to thank Eric Babson for serving on my thesis and qualifying exam committee and Anne Schilling and Nina Amenta for serving on my qualifying exam committee. Also, I am extremely grateful for the people who have been a part of the ADM, CCACAOO and MADDD seminars at Davis over the last several years and for all the wonderful talks I have been able to see.

Outside of the Davis, I am particularly grateful for the math events I attended at MSRI and UCB. In particular, I would like to thank Adam Sheffer and Joshua Zahl for hosting the Polynomial Method summer school at MSRI during the summer of 2019.

I also want to thank Luis Rademacher (again) and the National Science Foundation for the financial support I have received. This work was supported by the National Science Foundation under Grants CCF-2006994, CCF-1657939, CCF-1422830 and CCF-1934568.

Finally, I would like to thank my friends and family; for their support and love, for the lovely things we've seen together, and for asking me questions about math, and even occasionally asking good questions and being truly interested in the answers I gave.


CHAPTER 1

Introduction

This dissertation is about dividing finite sets of points into two equal parts. The main goal is to understand, in various situations, the number of distinct ways that a set of points can be divided into two equal parts. In order to better understand this sort of problem, we ask some new questions which are relevant to discrete geometry, convex geometry, and real algebraic geometry and which go beyond the basic question of equal division of finite point sets.

The ideas presented here begin with *halving lines*. Given a set of points in the plane, a *halving line* is a line that divides the set of points into two equal parts. The *halving lines problem* asks: What is the maximum number of distinct halving lines that a set of n points can have as a function of n ? The halving lines problem is now usually referred to as the *k-set problem* where a *k-set* of a set of n points in the plane is a subset of size k which can be separated from the remaining points by a line. The *k-set problem* asks one to determine the maximum number of *k-sets* that a set of n points can have. Of course, the halving lines problem is the $k = n/2$ case of the more general *k-set* problem.

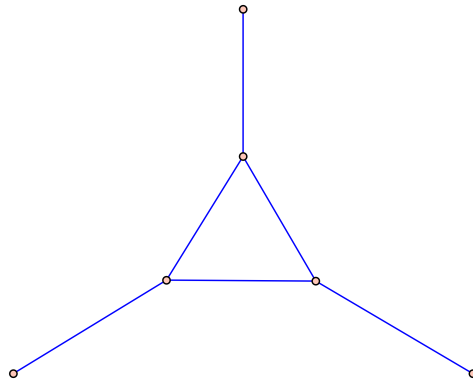
FIGURE 1.1. A set of 6 points with all halving lines drawn. This configuration of points has the maximum number of halving lines over all configurations of size 6.



The question of determining the maximum number of k -sets for point sets in the plane was first raised by A. Simmons in unpublished work. Straus, also in unpublished work, gave a construction showing an $\Omega(n \log n)$ lower bound for the $k = n/2$ case. Lovász [Lov71] published the first paper on k -sets, establishing an $O(n^{3/2})$ upper bound. See also the paper [ELSS73] of Erdős, Lovász, Simmons, and Straus. The main challenge is that even for the basic k -set problem on the plane, the asymptotics of the maximum number of k -sets is not well understood despite decades of effort. See Chapter 3 for an overview of the best known bounds.

In the late 1960s and early 1970s when the k -set problem was first studied, the field of computational geometry barely existed. Today, however, computational geometry is one of the main reasons for studying k -sets because they have important applications to geometric algorithms. For example, there are strong connections to order- k Voronoi diagrams [Der82], halfspace range search [CP86], convex hulls and k -hulls [CSY87].

FIGURE 1.2. A set of 6 points with all 2-edges drawn. (see Definition 2.2.4).



Many of the results proven here are about generalizations of or variations on the original halving lines problem: *Given a family of curves or (hyper)surfaces, and a set of points in \mathbb{R}^d , the aim is to understand how many combinatorially distinct ways there are to divide the set of points into two equal parts using one of the surfaces in the given family.* A generalization of this question asks, for any fixed integer k , how many ways there are to separate k of the points from the remaining points. The subsets of size k which can be separated from the remaining points by one of the surfaces in the chosen family are called \mathcal{F} - k -sets where the notation \mathcal{F} specifies the family of surfaces, i.e., the *set system*, see Chapter 2. The original k -set problem is the case when the family of surfaces is the

family of all lines in the plane. As in the case of the original k -set problem, the most interesting and difficult question in this context is how many \mathcal{F} - k -sets can there be? More precisely, for any integer n , what is the maximum number of \mathcal{F} - k -sets that a set of n points can have as a function of n and k ? For most families of surfaces, the halving case, i.e., the case when k is equal to $n/2$, is the most important case to consider. And if one can solve the problem in this case, the problem has essentially been solved. However, most of our results will be stated for arbitrary values of k , as the generalization is usually not difficult to obtain.

As outlined in the rest of this introduction, our quest to better understand the halving lines problem and its variations leads us to use diverse techniques from convex and discrete geometry as well as real algebraic geometry. Our use of these techniques leads us to ask, and in some cases answer, several new questions about algebraic or convex sets in \mathbb{R}^d . In this way, all of the questions/results in this dissertation are either about or inspired by k -sets.

Below is an outline of the content of each chapter of this document. Chapters 3 and 4 are based on the paper [LR20] written jointly with Luis Rademacher. Chapters 5 and 6 are based on the paper [LR21] which was also written jointly with Luis Rademacher.

Overview of Chapter 2: Preliminaries

This chapter reviews some of the most important existing definitions and theorems that are used throughout this dissertation. The definitions and theorems that we review come from the fields of discrete and convex geometry, combinatorics and probability, classical algebraic geometry, and mass partition problems.

Overview of Chapter 3: The k -set problem for algebraic set systems

This chapter initiates the study of the k -set problem for more general set systems. A *set system* is a pair (X, \mathcal{R}) , where X is a set (called the *universe*) and \mathcal{R} is a family of subsets of X . The subsets are called *ranges*. For the original k -set problem in the plane, the universe is \mathbb{R}^2 and the family of ranges is all halfplanes. We focus on *algebraic set systems*, i.e., set systems whose universe is \mathbb{R}^d and whose ranges are described by polynomial inequalities. This means that the separating surfaces are algebraic surfaces. Thus, the problems studied in this section are about the number of ways that one can separate k points from a given set of n points using an algebraic surface from a

chosen family. One of our most surprising results is when the family of surfaces is all conic sections in the plane. In this case, we prove a formula which exactly counts, given any set of n points in general position and any k , the number of conic sections which separate k of the points from the remaining points (Theorem 3.3.8). We also study the k -set problem for other algebraic set systems, and when we cannot count k -sets exactly, we at least establish upper bounds.

Overview of Chapter 4: Generally k -neighborly embedded manifolds and varieties

The proof of the result mentioned above about conic sections uses a very surprising property of the degree two Veronese map of the plane (Definition 3.2.2). This map has the property that it maps every generic set of points to a neighborly set of points (see Definition 2.1.2 for the definition of neighborly). We call embeddings with this property *generally neighborly embeddings* or *generally k -neighborly embeddings* for certain k (Definition 4.2.2). The image of a generally k -neighborly embedding of \mathbb{R}^d is called a generally k -neighborly d -manifold.

Recall that the *moment curve* M is the image of the map $\varphi : \mathbb{R} \mapsto \mathbb{R}^p$ where $\varphi(t) = (t, t^2, \dots, t^p)$. Generally neighborly d -manifolds should be seen as higher dimensional generalizations of the moment curve for the following reason. The most important property of the moment curve is that any configuration of points on M is neighborly. To find a higher-dimensional version of the moment curve, one would want to find some map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ (with $d \geq 2$) with the property that every set of points on $\varphi(\mathbb{R}^d)$ is neighborly. However, it was shown in [KW08] that no such map exists. Therefore, the best that we can ask for is that every *generic* set of points on $\varphi(\mathbb{R}^d)$ is neighborly. When this is the case, $\varphi(\mathbb{R}^d)$ is called a generally neighborly d -manifold. When $d = 2$, generally neighborly d -manifolds do exist. An example is the image of the degree 2 Veronese map of the plane and more examples are given in Chapter 4. However, we conjecture in Conjecture 4.2.5 that for $d \geq 3$, generally neighborly d -manifolds do not exist.

This conjecture is part of the following more general question which we leave open: What is the minimal dimension of the ambient space in which a generally k -neighborly d -dimensional manifold can exist? This question appears as Problem 4.3.2 in Chapter 4.

Because we are unable to resolve the main question we have about generally k -neighborly manifolds, we study a closely related question about *generally k -neighborly algebraic varieties* (Definition 4.4.1). Again, the main question is what is the minimal dimension of the ambient space in which a generally k -neighborly d -dimensional algebraic variety can exist? We show that the minimal dimension is $2k + d - 1$ (Theorem 4.4.3). The proof uses another neighborliness property called *weakly neighborly* (see Definition 4.5.1). It turns out that all generally k -neighborly varieties and manifolds are also weakly k -neighborly and using this property makes it easier to prove Theorem 4.4.3.

The questions in this section are in part inspired by a question asked by Micha Perles in 1982 and studied by Kalai and Wigderson [KW08]. Perles' question is about *neighborly embeddings* which we briefly discuss at the end of Chapter 4.

Overview of Chapter 5: Improved bounds for the expected number of k -sets

Chapter 5 returns to the original k -set problem in the plane, except that we study the random version of the problem. Instead of trying to determine the maximum number of k -sets that a set of n points can have, we study the expected number of k -sets of a set of n points chosen from some probability distribution on the plane. Our main contribution is an improved upper bound on the expected number of k -sets when the distribution is any probability distribution on the plane such that the measure of every line is 0. In this case, we show that the expected number of k -sets is $O(n^{5/4})$ (Theorem 5.1.2). The assumption that every line has measure 0 is a very minor restriction on the distribution. This is the first result of this type for probability distributions at this level of generality. Our result is interesting because our $O(n^{5/4})$ bound is significantly better than the current best known bound on the maximum number of k -sets for deterministic sets of points, which is $O(n^{4/3})$ [Dey97]. Furthermore, we have some reason to believe that the random version of the k -set problem that we study may be more or less equivalent to the original k -set problem. See Section 5.1 for a discussion of why we believe this may be true.

The proof of our bound on the expected number of k -edges begins by using vertical lines to partition the plane into open vertical strips of equal probability. We investigate how it may be possible to improve our bound by partitioning the plane using the continuous version of the

polynomial partitioning theorem of Guth and Katz rather than the basic partition by vertical lines. This motivates a question about the points of intersection of an algebraic curve and the k -edge graph of a set of points (Question 5.3.1).

Overview of Chapter 6: Translations of a fixed convex body in the plane

In this chapter we study the k -set problem for set systems for which the set of ranges consists of all translations of some strictly convex body. That is, for a convex body $C \subset \mathbb{R}^2$ and a set S of n points, we define a T_C - k -set of S to be a subset T of S of size k such that there exists a translation of C which contains T in its interior and contains no other points from S . For any strictly convex body C we determine bounds for the expected number of T_C - k -sets which are tight up to logarithmic factors (Theorem 6.2.9 and Theorem 6.5.3). The lower bound uses the uniform convergence theorem of Vapnik and Chervonenkis [VC71]. And the upper bound uses a technique due to Bárány and Steiger [BS94].

CHAPTER 2

Preliminaries

The point of this chapter is to collect some necessary definitions as well as state some theorems which hopefully will give the reader an idea of the types of techniques that we use.

2.1. Discrete geometry and convex geometry

Of all the classic results in discrete geometry, the most important one for our considerations is *Radon's theorem*

Theorem 2.1.1 (Radon's theorem). *Any set of $d + 2$ points in \mathbb{R}^d can be partitioned into two sets A, B so that the convex hulls of A and B have a common point.*

In fact, the theorem we use is a stronger version of Radon's theorem (Lemma 4.5.4) for the case when the $d + 2$ points are in general position.

Radon's theorem is relevant to Chapter 4 because of its connection to *neighborly* point configurations and polytopes.

By a *polytope*, we mean a *convex polytope*. See [Zie95] for background on polytopes.

Definition 2.1.2. *A polytope is k -neighborly if any set of k or fewer vertices forms a face. A d -dimensional polytope is neighborly if it is $\lfloor d/2 \rfloor$ -neighborly.*

If we are talking about point sets instead of polytopes, we will say that a set of points is *k -neighborly* (respectively, *neighborly*) if it is the vertex set of a k -neighborly (respectively, neighborly) polytope.

One way of producing neighborly point sets is by choosing a finite set of points on the *moment curve* $M := \{(t, t^2, \dots, t^d) : t \in \mathbb{R}\} \subset \mathbb{R}^d$. The moment curve is the standard example of an *order d curve* which is a curve that is intersected by any hyperplane in at most d points. Any finite set of distinct points on an order d curve is neighborly [MSRB71, Stu87].

The connection between Radon's theorem and neighborliness is that it is a simple consequence of Radon's theorem that if a d -dimensional polytope is $(\lfloor d/2 \rfloor + 1)$ -neighborly then it is a simplex. Our proof of a crucial lemma in Chapter 4 (Lemma 4.6.2) uses a very similar idea, just in a slightly more technical context.

From the area of convex geometry, a classic and important result we will use is about separation of convex sets in \mathbb{R}^d .

Two sets $Q, R \subseteq \mathbb{R}^p$ can be *weakly separated* if there exist a non-zero $a \in \mathbb{R}^p$ and $t \in \mathbb{R}$ such that $Q \subseteq \{x \in \mathbb{R}^p : a \cdot x \leq t\}$ and $R \subseteq \{x \in \mathbb{R}^p : a \cdot x \geq t\}$. We say that the hyperplane $\{x \in \mathbb{R}^p : a \cdot x = t\}$ *weakly separates* Q from R . This separation is said to be *proper* if Q and R are not both contained in $\{x \in \mathbb{R}^p : a \cdot x = t\}$.

Now we can state the so-called *separating hyperplane theorem*:

Theorem 2.1.3 (Theorem 1.3.8 in [Sch14]). *Let $Q, R \subset \mathbb{R}^d$ be non-empty convex sets. Then Q and R can be properly separated if and only if*

$$\operatorname{relint} Q \cap \operatorname{relint} R = \emptyset.$$

Another important notion of separation is called *strict separation*.

Definition 2.1.4. *Two sets $Q, R \subseteq \mathbb{R}^p$ can be strictly separated if there exist a non-zero $a \in \mathbb{R}^p$ and $t \in \mathbb{R}$ such that $Q \subseteq \{x \in \mathbb{R}^p : a \cdot x < t\}$ and $R \subseteq \{x \in \mathbb{R}^p : a \cdot x > t\}$. We say that the hyperplane $\{x \in \mathbb{R}^p : a \cdot x = t\}$ strictly separates Q from R .*

2.2. k -sets and k -facets

Definition 2.2.1. *Let S be a set of points in \mathbb{R}^d . A k -set of S is a subset $A \subset S$ of size k that can be strictly separated from $S \setminus A$ by a hyperplane.*

Definition 2.2.2. *We use $a_k(S)$ to denote the number of k -sets of the set $S \subset \mathbb{R}^d$ and $a_d(k, n)$ for the maximum number of k -sets that a set of n points in \mathbb{R}^d can have.*

When studying the k -set problem, one usually only considers point sets which are in general linear position.

Definition 2.2.3. A set of at least $d + 1$ points in \mathbb{R}^d is in general linear position if no $d + 1$ (and thus, fewer) points are affinely dependent.

This reduction is justified by the observation that the maximum number of k -sets is attained by a set of points in general linear position (see for example [Wag08]). For point sets in general linear position, one can study the closely related concept of k -facets.

Definition 2.2.4. Let S be a finite set of points in general linear position in \mathbb{R}^d and let Δ be a subset of d points from S . The subset Δ along with some orientation of the hyperplane $\text{aff } \Delta$ is a k -facet of S if the open halfspace on the positive side of $\text{aff } \Delta$ contains exactly k points from S . In \mathbb{R}^2 , k -facets are also known as k -edges.

Definition 2.2.5. We use $e_k(S)$ to denote the number of k -facets of the set $S \subset \mathbb{R}^d$ and $e_d(k, n)$ for the maximum number of k -facets that a set of n points in general linear position in \mathbb{R}^d can have.

It seems unlikely that one would be able to determine $e_d(k, n)$ or $a_d(k, n)$ precisely, so instead efforts have focused on finding the asymptotic behavior of these functions.

The k -set problem asks one to determine the asymptotic behavior of $a_d(k, n)$. And the k -facet problem asks one to determine the asymptotic behavior of $e_d(k, n)$.

If one is only concerned with the asymptotics, then it suffices to study either k -sets or k -facets since for fixed d and $n \rightarrow \infty$, $a_d(k, n)$ and $e_d(k, n)$ have the same asymptotic behavior [Wag08].

2.3. Mass partitions

Mass partition theorems describe how a given set of points or collection of measures on \mathbb{R}^d is partitioned after dividing \mathbb{R}^d into a number of disjoint (usually open) sets.

Probably the most famous theorem of this sort is the *ham sandwich theorem*:

Theorem 2.3.1 (Ham sandwich theorem [ST42]). Given d absolutely continuous¹ measures μ_1, \dots, μ_d on \mathbb{R}^d , there exists a hyperplane H so that $\mu_i(H^+) = \mu_i(H^-)$ for all $1 \leq i \leq d$. (H^+ denotes the open halfspace on the positive side of H and H^- is the open halfspace on the negative side of H .)

¹The absolutely continuous measures are precisely those which have a density.

Part of the reason this theorem is interesting is that its proof is via the Borsuk-Ulam theorem from algebraic topology.

The ham sandwich theorem is relevant to our considerations because it is how one proves *polynomial partitioning*-type theorems which we make use of in Chapter 5.

The polynomial partitioning theorem of [GK15] has recently been used to solve a number of problems in discrete and combinatorial geometry [Gut16a]. It has also been used to give alternative proofs of some known results, see [KMS12]. Perhaps the most commonly used version of the polynomial partitioning theorem is the following, which we refer to as the discrete version;

Theorem 2.3.2 (Discrete polynomial partitioning [GK15]). *Let $S \subset \mathbb{R}^d$ be a set of n points. Then for each $r \leq n$ there is a non-zero polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree $O_d(r)$ such that $\mathbb{R}^d \setminus Z(f)$ is the union of a family \mathcal{O} of r^d pairwise disjoint open sets such that each $O \in \mathcal{O}$ contains at most n/r^d points of S . (The notation O_d means that the constants in the bounds depend only on d .)*

In Chapter 5, we are focused on bounding the expected number of k -edges of a sample of points from a distribution on the plane (Theorem 5.3.2). The idea of the proof of this theorem is to use a divide-and-conquer approach which first partitions the plane into a number of cells of equal probability. We then consider separately the expected number of k -edges which cross from one cell to another and the expected number of k -edges that do not cross the boundary of the partition. Therefore, we need a partitioning theorem which applies to probability distributions rather than finite point sets. This type of result, which we refer to as the continuous version of the polynomial partitioning theorem, has been used to establish improved bounds for the restriction problem in harmonic analysis, see [Gut16b] as a starting point.

Theorem 2.3.3 (Continuous polynomial partitioning [Gut16b]). *Let $W \in L^1(\mathbb{R}^d)$ with $W \geq 0$. Then for each r , there is a non-zero polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree at most r such that $\mathbb{R}^d \setminus Z(f)$ is the union of a family \mathcal{O} of $\Theta_d(r^d)$ pairwise disjoint open sets such that for all $O \in \mathcal{O}$, the integrals $\int_O W$ are equal.*

The open sets O in the above theorems are called the *cells* of the partition.

Using the density of the non-singular polynomials in the space of all polynomials of fixed degree in d variables (see Figs. 2.1a and 2.1b), it is possible to obtain, as a corollary of the continuous




FIGURE 2.1. An example of a singular degree 3 algebraic curve and a non-singular degree 3 algebraic curve which are “close” in the space of all degree 3 algebraic curves.

polynomial partitioning theorem, a version where all the irreducible components of the dividing surface $Z(f)$ are non-singular varieties.

Theorem 2.3.4 (Non-singular continuous polynomial partitioning [Gut16b]). *Let $W \in L^1(\mathbb{R}^d)$ with $W \geq 0$. Then for each r , there is a non-zero polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ of degree at most r such that $\mathbb{R}^d \setminus Z(f)$ is the union of a family \mathcal{O} of $\Theta_d(r^d)$ pairwise disjoint open sets such that for all $O \in \mathcal{O}$, the integrals $\int_O W$ are within a factor of two of each other. Furthermore, all irreducible components of $Z(f)$ are non-singular.*

Notice that the above theorems do not allow us to partition an arbitrary probability distribution, they only apply to those for which there exists a density W . This is the reason why our proof in Section 5.3 only applies to distributions which have a density. As far as we know, it is an open problem to prove a polynomial partitioning theorem similar to Theorem 2.3.3 for arbitrary probability distributions. The proof of Theorem 2.3.3 relies on the ham sandwich theorem stated above which only applies to distributions which have a density. In order to extend the polynomial partitioning theorem to arbitrary measures, one could attempt to use the more general form of the ham sandwich theorem and repeat the proof of Theorem 2.3.3. However, since an arbitrary distribution may be concentrated in a hyperplane or even a point, one would not be able to guarantee that the integrals $\int_{O_i} W$ are all equal, but only that all integrals are at most some fixed value.

2.4. Set systems and VC-theory

Definition 2.4.1. A set system is a pair (X, \mathcal{R}) , where X is a set (called the universe) and \mathcal{R} is a family of subsets of X . The elements of \mathcal{R} are called ranges).

We will use the language of set systems in Chapter 3 when we define k -sets for set systems other than halfspaces.

Set systems are also the starting point for VC-theory. The idea of VC-theory is to come up with a way to measure the complexity of a geometric object or family of geometric objects.

Our main use of VC-theory will be of the uniform convergence theorem of Vapnik and Chervonenkis [VC71]. We state this theorem below as Theorem 2.4.6. The idea of this theorem is as follows. Say (X, \mathcal{R}) is some set system and P is some probability distribution on X . For the uniform convergence theorem to hold, we also need to assume that the VC-dimension (Definition 2.4.2) of the set system (X, \mathcal{R}) is finite. Let X_1, \dots, X_m be a sample of m iid points from P . For each $R \in \mathcal{R}$ the *relative frequency* of R with respect to the sample X_1, \dots, X_m is the quotient $\frac{|\{X_i: X_i \in R\}|}{m}$. The uniform convergence theorem then says that the relative frequencies of the events R converge uniformly to their probabilities. Another way of saying this is that for any $\epsilon > 0$, the probability that the maximum difference over all $R \in \mathcal{R}$ of the relative frequency of R and the probability of R is less than ϵ converges to 0 as the size of the sample tends to infinity.

First we define *VC-dimension*. For a set system (X, \mathcal{R}) and a subset $Y \subset X$, we say that a subset $A \subseteq Y$ is *induced* by a range $R \in \mathcal{R}$ if $A = Y \cap R$.

For a subset $Y \subset X$, we say that Y is *shattered* by (X, \mathcal{R}) if $|\{R \cap Y : R \in \mathcal{R}\}| = 2^{|Y|}$. In other words, Y is shattered by the set system (X, \mathcal{R}) if all subsets of Y can be induced by intersecting Y with one of the ranges in \mathcal{R} .

Definition 2.4.2. The VC-dimension of set system (X, \mathcal{R}) is the maximum size of a subset of X that is shattered by (X, \mathcal{R}) .

Example 2.4.3 (Halfspaces). Consider the set system (X, \mathcal{R}) where $X = \mathbb{R}^d$ and \mathcal{R} is all closed halfspaces. The VC-dimension of this set system is $d + 1$. This follows from Radon's theorem.

Definition 2.4.4 ([VC71]). The growth function $\pi_{\mathcal{R}}(n)$ of set system (X, \mathcal{R}) is the maximum number of subsets of a set $Y \subseteq X$ of size n that can be induced by ranges in \mathcal{R} . So the growth

function is given by

$$\pi_{\mathcal{R}}(n) = \max_{A \subseteq X, |A|=n} |\{A \cap R : R \in \mathcal{R}\}|.$$

Theorem 2.4.5 ([VC71]). *The VC-dimension gives a bound on the growth function: If the VC-dimension of the set system (X, \mathcal{R}) is d , then*

$$\pi_{\mathcal{R}}(n) \leq \sum_{i=0}^d \binom{n}{i}.$$

Theorem 2.4.6 (uniform convergence [VC71]). *Suppose that (X, \mathcal{R}) is a set system and P is a probability distribution on X such that every $R \in \mathcal{R}$ is measurable according to P and also so that the function $\sup_{R \in \mathcal{R}}(\cdot)$ as defined below is a random variable (i.e., is measurable). For any $0 < \epsilon < 1$, and $\{X_1, \dots, X_m\}$ a sample of $m \geq 2/\epsilon^2$ points from P ,*

$$\mathbb{P} \left(\sup_{R \in \mathcal{R}} \left| P(R) - \frac{|\{X_i : X_i \in R\}|}{m} \right| < \epsilon \right) \geq 1 - 4\pi_{\mathcal{R}}(2m)e^{-\epsilon^2 m/8}.$$

In Chapter 6 we will also need to use the notion of the *dual set system*.

Definition 2.4.7 (dual set system). *The dual set system of set system (X, \mathcal{R}) is the set system (\mathcal{R}, X^*) where X^* is the family of sets of the form $\{R \in \mathcal{R} : x \in R\}$ for $x \in X$.*

The *dual growth function* of (X, \mathcal{R}) is the growth function of its dual set system.

2.5. Classical real algebraic geometry

Given polynomials $f_1, f_2, \dots, f_r \in \mathbb{R}[x_1, \dots, x_p]$, the *affine algebraic variety* defined by the f_i 's is the set

$$V(\mathbb{C}) = \{x \in \mathbb{C}^p : f_1(x) = 0, \dots, f_r(x) = 0\}.$$

Algebraic varieties are most well behaved when the field k is algebraically closed, for example $k = \mathbb{C}$. Our applications, however, require us to study algebraic varieties over the real numbers. In other words, for a given variety $V(\mathbb{C}) \subset \mathbb{C}^p$, we are mainly interested in the subset $V(\mathbb{R}) := V(\mathbb{C}) \cap \mathbb{R}^p$ of real points. We refer to $V(\mathbb{R})$ as an *affine real algebraic variety*. Unless otherwise stated, by an *algebraic variety* or simply *variety* we mean an affine real algebraic variety.

Since we are mainly interested in the set of real points, we will write V for $V(\mathbb{R})$ and will write $V(\mathbb{C})$ to indicate when the complex points are also considered. See [BCR98, Har95] for more definitions from real algebraic geometry.

Definition 2.5.1. *A variety $V \subset \mathbb{R}^p$ (resp. $V(\mathbb{C}) \subset \mathbb{C}^p$) is non-degenerate if it is not contained in any hyperplane in \mathbb{R}^p (resp. \mathbb{C}^p).*

One of the most useful properties of varieties is that there is simple and canonical way to decompose them into basic components.

Definition 2.5.2. *An algebraic variety is irreducible if it cannot be written as the union of two proper algebraic subvarieties.*

Any variety V has a unique decomposition into irreducible components. That is, there is a unique way to write $V = V_1 \cup \dots \cup V_m$ where each V_i is an irreducible subvariety and no V_i is contained in any V_j for $j \neq i$. In such a decomposition, the V_i are called the *irreducible components* of V .

We will use facts about the smooth and singular points of a variety. Suppose $V(\mathbb{C})$ is a variety and the ideal of $V(\mathbb{C})$ is generated by the polynomials f_1, \dots, f_r . The *smooth* points of $V(\mathbb{C})$ are those points where the Jacobian matrix of the f_i 's has maximal rank. A *singular point* of a variety is a point that is not smooth.

Definition 2.5.3. *We use V_{sm} to denote the set of smooth points of an algebraic variety V , V_{sing} is the set of singular points.*

Given a real variety $V \subset \mathbb{R}^p$, let $V_{\mathbb{C}}$ denote the smallest complex variety which contains V . It is well known that $V_{\mathbb{C}}$ is unique and furthermore that there is a bijection between irreducible components of V and irreducible components of $V_{\mathbb{C}}$, see [Whi57]. Note that $V(\mathbb{C})$ is not always equal to $V_{\mathbb{C}}$. However, it is a useful fact that whenever V contains a smooth real point, V is Zariski

dense in $V(\mathbb{C})$ and so $V(\mathbb{C}) = V_{\mathbb{C}}$, see [BCR98, Section 2.8].² We will always assume that our varieties contain smooth real points.

By the *dimension* of a real algebraic variety V we will mean the dimension of $V(\mathbb{C})$. There is another notion of dimension for real algebraic varieties.

Definition 2.5.4. *The real dimension of a real algebraic variety V is the maximal integer d such that there is a homeomorphism of $[0, 1]^d$ into some subset of V .*

The real dimension of V does not always equal the dimension of V . However, the real dimension of V is never more than the dimension of V (see [BCR98, Proposition 2.8.14]) and if V contains a smooth real point, then these dimensions do agree. This is because around any smooth real point of a d -dimensional variety $V \subset \mathbb{R}^p$, there is a neighborhood which is a smooth d -dimensional submanifold of \mathbb{R}^p [BCR98, Proposition 3.3.11].

In one of the proofs in Section 4.6 we will also consider projective varieties, see [Har95] for background on projective algebraic geometry. We use $\mathbb{P}^p(K)$ to denote p -dimensional projective space over the field $K = \mathbb{C}$ or $K = \mathbb{R}$.

²There is a potential confusion here about the existence of smooth points in real varieties. When working with varieties, if you start with a variety $V \subset \mathbb{R}^d$ and then take V to be defined by the polynomials $I(V)$, then one can prove that every real variety contains smooth real points and thus $V(\mathbb{C}) = V_{\mathbb{C}}$ always. This is not the approach we take. (Recall that we defined real varieties as the set of real points of a complex variety.) For this reason, in our results later on about real varieties we always add the assumption that they contain smooth real points. But this assumption is only necessary because of the way we defined real varieties.

CHAPTER 3

The k -set problem for algebraic set systems

In this chapter we study a variation on the k -set/ k -facet problem with hyperplanes replaced by algebraic surfaces. In stark contrast to the original k -set/ k -facet problem, there are some natural families of algebraic curves for which the number of k -facets can be counted exactly. For example, we show that the number of halving conic sections for any set of $2n + 5$ points in general position in the plane is $2\binom{n+2}{2}^2$. Additionally, we give a simple argument which improves the best known bound on the number of k -sets/ k -facets for point sets in convex position.

3.1. Introduction

It is natural to ask questions similar to the k -set problem but for families of surfaces different from all hyperplanes. These sorts of questions have been studied in [Der82, Ard04, BHP08, CFSS14, CFSS20]. Ardila's paper [Ard04] shows that for any set of $2n + 1$ points in general position in the plane and any $0 \leq k \leq 2n - 2$, the number of circles that go through 3 points and have k points on one side is exactly $2(k + 1)(2n - k - 1)$. We call this phenomenon *exact counting*: when for a family of curves (or surfaces), there exists an integer d such that given any generic set of n points and any k , the number of curves which pass through d points and have k points on one side depends only on n and k and not on the points. A result essentially equivalent to Ardila's was proven earlier in [Der82] by counting vertices of certain Voronoi diagrams. Chevallier et al. extended the result to convex pseudo-circles in [CFSS20].

Borrowing the language of set theory/computational geometry/learning theory, one can think of the k -set problem as being formulated over a *set system* (also known as a *hypergraph*, *hypothesis class* or *range space*), namely a universe and a family of subsets of the universe. In the k -set problem the universe is \mathbb{R}^d and the family of subsets is all halfspaces. This paper takes a step towards the understanding of the k -set problem for general set systems. We focus on set systems induced by maps in the following way: given a map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$, the set system induced by φ

has universe \mathbb{R}^d and family of subsets $\{\varphi^{-1}(H) : H \text{ is a closed halfspace in } \mathbb{R}^p\}$. Moreover, most of our results involve maps φ with components that are polynomials, so that the separating surfaces in the resulting set system are algebraic surfaces. One of our main examples is the Veronese map (Definition 3.2.2) which induces separators that are algebraic surfaces of degree at most m . The Veronese map is also known as the feature map of the polynomial kernel in machine learning [STC04].

Our contributions:

- **Exact count.** We show that the exact count phenomenon of [Der82, Ard04] (for halving circles) holds for other natural set systems: conic sections (Theorem 3.3.8) and homogeneous polynomials of fixed even degree on the plane (Theorem 3.3.9). We prove this by establishing a remarkable property of the corresponding maps: generic point sets are mapped to point sets that form the vertices of a neighborly polytope (Theorem 3.3.4 and Theorem 3.3.5, see Section 3.2.5 for background). This is then combined with the known fact that the number of k -facets of a neighborly point set is given by a formula that depends only on the dimension, on k , and on the number of points [Cla87, CS89], [Wag08, Proposition 4.1].
- **Convex position bound.** We show an improved upper bound on the number of k -sets/ k -facets for points in convex position (Theorem 3.3.12). While our argument is simple, we are not aware of any known bounds for the convex case better than the general case in dimension higher than three (the convex case is well understood in two and three dimensions).
- **Degree of neighborliness.** We study the degree of neighborliness of point sets mapped by a φ with components “all monomials of degree at most m ” or “all monomials of degree exactly m ” (Theorem 3.3.10 and Theorem 3.3.11). In particular, for even m , point sets are mapped into point sets in convex position and the convex position bound gives an improved bound on the number of k -facets.

Outline of the Chapter. Section 3.2 reviews some preliminary material and introduces our generalization of the k -facet problem to set systems which are induced by maps. In Section 3.3 we

count k -facets for set systems induced by maps. This amounts to studying k -facets of point sets of the form $\varphi(S)$ where $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ is some map and $S \subset \mathbb{R}^d$ is a finite set of points.

3.2. Preliminaries

3.2.1. Generic properties and general position. After defining k -sets and k -facets for set systems other than halfspaces, we will need to use various notions of general position different from general linear position.

A *generic property* (of point sets) is one that holds for all but a relatively small number of atypical point sets. The point sets which satisfy a generic property are said to be *in general position*. We use the terms “generic point set” and “point set in general position” interchangeably.

In algebraic geometry a generic property is one that holds for a dense and open set. In other fields, a generic property holds almost everywhere. For concreteness in some of our statements we set “generic property” to mean a property that holds in an open and dense set, but this choice is not always crucial.

Now we state more explicitly what it means for a point configuration to be generic. The collection of all configurations of n -point sets in \mathbb{R}^d can be identified with \mathbb{R}^{dn} . If $G_n \subset \mathbb{R}^{dn}$ is the collection of all generic point configurations of size n , then G_n should be dense and open in \mathbb{R}^{dn} for all n . See [Mat02, Section 1.1] for more on the meaning of general position in discrete geometry.

3.2.2. Set systems induced by maps. Recall that a set system is a pair (X, \mathcal{F}) where X is a ground set (or universe) and \mathcal{F} is a collection of subsets of X .

We will restrict our attention to set systems which are *induced by maps*. Suppose we have a map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$, that is, a map of \mathbb{R}^d into some (usually higher dimensional) space. Any such map induces a set system on the ground set \mathbb{R}^d in the following way. Let \mathcal{F}_φ consist of all regions $R \subset \mathbb{R}^d$ of the form $\varphi^{-1}(H)$ where H is a closed halfspace in \mathbb{R}^p . We say that R is induced by the halfspace H and we say that the set system $(\mathbb{R}^d, \mathcal{F}_\varphi)$ is induced by φ . Many interesting set systems are induced by maps.

3.2.3. Set systems induced by Veronese-type maps.

Definition 3.2.1. A polynomial map is a map $\mathbb{R}^d \rightarrow \mathbb{R}^p$ defined by

$$x \mapsto (f_1(x), \dots, f_p(x))$$

where f_1, \dots, f_p are polynomials.

Here we introduce our primary examples of polynomial maps and set systems.

Definition 3.2.2. The degree m Veronese map of \mathbb{R}^d is the map $V_m^d : \mathbb{R}^d \rightarrow \mathbb{R}^{\binom{d+m}{m}-1}$ which maps (x_1, \dots, x_d) to the vector (in some order) of all non-constant monomials of degree at most m in the d variables x_1, \dots, x_d .

Definition 3.2.3. The degree m homogeneous Veronese map of \mathbb{R}^d is $HV_m^d : \mathbb{R}^d \rightarrow \mathbb{R}^{\binom{d+m-1}{m}}$ which maps (x_1, \dots, x_d) to the vector (in some order) of all monomials of degree m in the d variables x_1, \dots, x_d .

We will use the notation $(\mathbb{R}^d, \mathcal{P}_m^d)$ for the set system induced by the degree m Veronese map of \mathbb{R}^d and $(\mathbb{R}^d, \mathcal{H}_m^d)$ for the set system induced by the degree m homogeneous Veronese map of \mathbb{R}^d .

As a concrete example, consider the degree 2 Veronese map of \mathbb{R}^2 . This is the map $V_2^2 : \mathbb{R}^2 \rightarrow \mathbb{R}^5$ where $V_2^2(x, y) = (x^2, xy, y^2, x, y)$. The set system induced by this map is $(\mathbb{R}^2, \mathcal{P}_2^2)$. Its subsets consist of all regions of the plane determined by some conic section.

3.2.4. On k -sets and k -facets for set systems induced by maps. The natural notion of a k -set of a set system is a range that contains exactly k points. It is often more convenient to work with k -facets, but it is not clear how to define them for set systems whose ranges lack a well-defined boundary. This motivates our restriction to set systems induced by maps, as their ranges have a well-defined boundary and interior.

Definition 3.2.4. Given a set system (X, \mathcal{F}_φ) induced by a map $\varphi : X \rightarrow \mathbb{R}^p$ and a finite set $S \subset X$, an \mathcal{F}_φ - k -set of S is a subset $A \subset S$ of size k such that $\varphi(A)$ can be strictly separated from $\varphi(S \setminus A)$ by a hyperplane.

Definition 3.2.5. Given a set system (X, \mathcal{F}_φ) induced by a map $\varphi : X \rightarrow \mathbb{R}^p$ and a finite set $S \subset X$ such that $\varphi(S)$ is in general linear position, an \mathcal{F}_φ - k -facet of S is a subset P of p points

from S , along with some orientation of the hyperplane $\text{aff } \varphi(P)$, such that the subset $\varphi(P)$ along with the chosen orientation of $\text{aff } \varphi(P)$ is a k -facet of $\varphi(S)$.

Observe that counting \mathcal{F}_φ - k -sets/facets simply amounts to counting k -sets/facets of point sets of the form $\varphi(S)$. Therefore, upper bounds for $e_d(k, n)$ and $a_d(k, n)$ immediately imply non-trivial upper bounds on the number of \mathcal{F}_φ - k -sets/facets that a set of points may have:

Proposition 3.2.6. *Given a set system (X, \mathcal{F}_φ) induced by a map $\varphi : X \rightarrow \mathbb{R}^p$, and a finite subset $S \subset X$ such that $\varphi(S)$ is in general linear position, the number of \mathcal{F}_φ - k -facets of S is at most $e_p(k, n)$.*

For \mathcal{F}_φ - k -sets, we do not need to assume that $\varphi(S)$ is in general linear position since k -sets are defined for any point set whether or not it is in general linear position.

Proposition 3.2.7. *Given a set system (X, \mathcal{F}_φ) induced by a map $\varphi : X \rightarrow \mathbb{R}^p$, the number of \mathcal{F}_φ - k -sets that a set of n points in X may have is at most $a_p(k, n)$.*

PROOF. In the case when φ is not injective, $\varphi(S)$ may need to be considered as a multiset. Therefore, we start with the observation that $a_p(k, n)$ is the maximum number of k -sets even for point sets which have repeated points, i.e., multisets of points. To see this, observe that perturbing a set of points can only increase the number of k -sets [Wag08]. Therefore, if we start out with a multiset, we can perturb it slightly to create a set (in general linear position) with the same number of points and at least as many k -sets. Now, the number of \mathcal{F}_φ - k -sets of S is equal to the number of k -sets of $\varphi(S)$, which is at most $a_p(k, n)$. \square

3.2.5. Neighborly polytopes. For a set of n points in convex position in the plane, the number of k -facets is precisely n for all values of k . In \mathbb{R}^3 , a similar result is true: the number of k -facets for a set S of n points in general position which form the vertex set of a 3-polytope is $2(k+1)n - 4\binom{k+2}{2}$, see [Wag08]. There is no such result in dimension $d \geq 4$, i.e., convex position does not force a point set in \mathbb{R}^d ($d \geq 4$) to have a specific number of k -facets. In fact, the k -set/ k -facet problem for point sets in convex position in \mathbb{R}^4 is only slightly better understood than the problem for arbitrary point sets, see Theorem 3.3.12.

However, if we assume that our point set is not only in convex position but is also *neighborly*, then $e_k^{(d)}(S)$ is determined precisely by $|S|$ and k . See Definition 2.1.2 for the definition of neighborly point sets and [Wag08, Zie95] for a more thorough introduction to neighborly polytopes.

Proposition 3.2.8 ([Cla87, CS89], [Wag08, Proposition 4.1]). *Let S be a neighborly set of n points in general linear position in \mathbb{R}^d . Then*

$$e_k(S) = \begin{cases} 2 \binom{k+\lceil d/2 \rceil - 1}{\lceil d/2 \rceil - 1} \binom{n-k-\lceil d/2 \rceil}{\lceil d/2 \rceil - 1} & \text{if } d \text{ is odd} \\ \binom{k+d/2-1}{d/2-1} \binom{n-k-d/2}{d/2} + \binom{k+d/2}{d/2} \binom{n-k-d/2-1}{d/2-1} & \text{if } d \text{ is even.} \end{cases}$$

3.3. Counting k -facets via maps

In this section we count \mathcal{F}_φ - k -facets when \mathcal{F}_φ is a set system induced by a map. When the map φ has certain properties, we can say more about the number of \mathcal{F}_φ - k -facets.

3.3.1. Counting k -facets exactly. It turns out that the maps associated to several families of polynomials we have discussed have the surprising property that they map generic point sets into the set of vertices of a neighborly polytope. Given such a map φ , we are able to exactly count the number of \mathcal{F}_φ - k -facets for point sets in general position.

Before stating the new results, we recall a result of [Der82, Ard04] which served as motivation. A *halving circle* of a point set of size $2n + 1$ is a circle which has 3 points on its boundary and $n - 1$ points on either side. In the following theorem general position means that no three points are collinear and no four are concyclic.


Theorem 3.3.1 ([Der82, Ard04]). *Any set S of $2n + 1$ points in general position in the plane has exactly n^2 halving circles. More generally, for any $0 \leq k \leq 2n - 2$, the number of circles that have 3 points of S on their boundary and k points on one side is exactly $2(k + 1)(2n - k - 1)^1$.*

The proof of Theorem 3.3.1 in [Ard04] is by a continuous motion argument. However, as noted there, it is possible to give a shorter proof using the method of maps as follows. The set system of all circles in the plane can be described as $(\mathbb{R}^2, \mathcal{F}_C)$ where $C : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is the map $C(x, y) = (x, y, x^2 + y^2)$. Since C maps generic point sets into convex position (on the surface of

¹This formula counts each halving circle twice, once for each orientation.

a paraboloid), Theorem 3.3.1 follows from an application of the formula in Proposition 3.2.8 since any 3-polytope is neighborly.

FIGURE 3.1. A set of 7 points with all *halving conic sections* drawn.



Now we define halving polynomials for other families of polynomials. Informally, for a finite set $S \subseteq \mathbb{R}^2$, a *halving conic section* of S is a conic section inequality having 5 points on its boundary and half of the remaining points of S in its interior. Unlike Theorem 3.3.1 on the circle problem above, we count halving conic sections twice, once for each orientation. This is to be consistent with the standard definition of k -facets (Definition 3.2.5). More precisely,

Definition 3.3.2. For a set S of $2n + 5$ points in \mathbb{R}^2 , a halving conic section is an \mathcal{F}_V - n -facet of S where $V := V_2^2$ is the degree 2 Veronese map of \mathbb{R}^2 .

Definition 3.3.3. For a set S of $2n + m + 1$ points in \mathbb{R}^2 , a halving homogeneous polynomial of degree m of S is an \mathcal{H}_m^2 - n -facet of S .

The halving case is a particular case of the more general problem of counting k -facets. It is generally believed that the maximum number of k -facets is maximized by the halving case, so it is considered the most important. However, we state our results counting \mathcal{H}_m^2 - k -facets and \mathcal{F}_V - k -facets for arbitrary values of k .

Theorem 3.3.4. *Assume a finite set of points $S \subseteq \mathbb{R}^2$ is in general linear position. Then the image of S by the degree 2 Veronese map V_2^2 is neighborly.*

PROOF. First we verify that $V_2^2(S)$ is the set of vertices of a polytope. There is a bijection between conic sections passing through points of S and hyperplanes passing through the images of those points by V_2^2 . Therefore, for every point $v \in S$ we need to find a conic section inequality passing through v and with all other points on one side. We can use an inequality of the form $(x - a)^2 + (y - b)^2 \leq r$ and adjust the constants a, b, r so that $(x - a)^2 + (y - b)^2 \leq r$ defines a circle that contains v on its boundary and has radius small enough so that no other point in S is in the circle. Now we show that $V_2^2(S)$ is neighborly. For every 2 points v_1, v_2 of S we need to find a conic section inequality passing through those points and with all other points on one side. One way to accomplish this is to use the line $ax + by = c$ through v_1 and v_2 . Then $(ax + by - c)^2 \leq 0$ is the required conic section inequality. \square

In terms of the terminology defined in Section 4.2, the above result says that V_2^2 is a “generally neighborly embedding”. The same result holds for the even degree homogeneous Veronese map of the plane.

Theorem 3.3.5. *Assume m is even and $S \subseteq \mathbb{R}^2$ is in general position, meaning that no two points of S lie on a common line through the origin. Then the image of S by HV_m^2 is neighborly.*

PROOF. The proof is similar to that of Theorem 3.3.4. For any set $\{v_1, \dots, v_k\}$ of $k \leq m/2$ points of S we need to find a degree m homogeneous polynomial inequality which passes through all the v_i and has all other points of S on one side. Let $v_{k+1}, \dots, v_{m/2}$ be points in the plane that belong to no line passing through the origin and a point of S . For each i , $1 \leq i \leq m/2$, let $a_i x + b_i y = 0$ be the line through the origin and v_i . Then $\prod_{i=1}^{m/2} (a_i x + b_i y)^2 \leq 0$ is a polynomial inequality with the required properties. \square

The next results require us to strengthen our general position assumptions from Theorem 3.3.4 and Theorem 3.3.5.

Definition 3.3.6. *A set $S \subset \mathbb{R}^2$ is in general position with respect to conics if S is in general linear position and $V_2^2(S)$ is in general linear position.*

Definition 3.3.7. A set $S \subset \mathbb{R}^2$ is in general position with respect to degree m homogeneous polynomials if no two points in S lie on a common line through the origin and $HV_m^2(S)$ is in general linear position.

Theorem 3.3.8. Any set of n points of \mathbb{R}^2 in general position with respect to conics has exactly $2\binom{k+2}{2}\binom{n-k-3}{2}$ \mathcal{F}_V - k -facets where $V := V_2^2$ is the degree 2 Veronese map of the plane.

PROOF. Let $S \subset \mathbb{R}^2$ be a set of n points in general position with respect to conics. There is a bijection between conic sections passing through 5 points of S and hyperplanes passing through 5 points of $V(S)$. Furthermore, there is a bijection between \mathcal{F}_V - k -facets of S and k -facets of $V(S)$. By Theorem 3.3.4, $V(S)$ is neighborly. Also, since S is in general position with respect to conics, $V(S)$ is in general linear position. Therefore the number of \mathcal{F}_V - k -facets of S is given by the formula from Proposition 3.2.8. \square

Theorem 3.3.9. Assume m is even. Any set of $2n + m + 1$ points of \mathbb{R}^2 in general position with respect to degree m homogeneous polynomials has exactly $2\binom{k+m/2}{m/2}\binom{n-k-m/2-1}{m/2}$ \mathcal{H}_m^2 - k -facets.

PROOF. Let $S \subset \mathbb{R}^2$ be a set of $2n + m + 1$ points in general position with respect to degree m homogeneous polynomials. As in the last proof, there is a bijection between \mathcal{H}_m^2 - k -facets of S and k -facets of $HV_m^2(S)$. By Theorem 3.3.5, $HV_m^2(S)$ is neighborly. Also, since S is in general position with respect to degree m homogeneous polynomials, $HV_m^2(S)$ is in general linear position. Since $HV_m^2 : \mathbb{R}^2 \rightarrow \mathbb{R}^{m+1}$, the formula from Proposition 3.2.8 in the case $d = m + 1$ completes the proof. \square

3.3.2. Lifting the moment curve. We say that a set system (X, \mathcal{F}) has the exact count property if the number of \mathcal{F} - k -sets for any set of n points in general position depends only on n and k and not on the configuration of points. We can generate many more set systems with the exact counting property by a lifting of the moment curve.

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function such that $\{(x, x') \in \mathbb{R}^{2d} : f(x) \neq f(x')\}$ is open and dense and let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be any function. We will say that a set $S = \{s_i\}_{i \in [n]}$ of n points in \mathbb{R}^d is in general position if $\prod_{i,j \in [n], i \neq j} (f(s_i) - f(s_j)) \neq 0$. Note that this is a reasonable definition of general position since if we are considering n -point sets, then point sets (in \mathbb{R}^{nd}) in general position

are open and dense in \mathbb{R}^{nd} . Assume that $m \geq 2$ is even. The map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^{m+1}$ given by

$$(3.1) \quad \varphi(x) = \left(f(x), (f(x))^2, \dots, (f(x))^m, g(x) \right)$$

satisfies that, for any set S of n points in general position in \mathbb{R}^d , $\varphi(S)$ is neighborly.² To see this, note that, since m is even, $\lfloor \frac{m+1}{2} \rfloor = m/2$. The projection of $\varphi(S)$ to the first m coordinates is $m/2$ -neighborly since it is a set of n distinct points on the moment curve in \mathbb{R}^m . We claim that this implies that $\varphi(S)$ is $m/2$ -neighborly as well: Let $\pi(\varphi(S))$ denote the projection to the first m coordinates. The points of $\varphi(S)$ all project to distinct vertices of $\pi(\varphi(S))$. By neighborliness of $\pi(\varphi(S))$, every subset F of at most $m/2$ vertices of $\pi(\varphi(S))$ forms a face. For any such face, there is a supporting hyperplane H . The preimage $\pi^{-1}(H)$ of H under the projection π is a hyperplane with normal having last coordinate 0. Moreover, $\pi^{-1}(H)$ is the supporting hyperplane for a face of $\varphi(S)$ formed by the lifted vertices $\pi^{-1}(F)$. This shows that $\varphi(S)$ is neighborly.

Given any admissible choice of functions f, g , the map above induces a set system $(\mathbb{R}^d, \mathcal{F}_\varphi)$ with the exact count property.

3.3.3. Improved bounds for \mathcal{P}_m^d - k -facets and \mathcal{H}_m^d - k -facets. Results like Theorem 3.3.8 and Theorem 3.3.9 are not possible for any of the other polynomial set systems we have discussed. However, some progress can be made.

Recall that in Proposition 3.2.6 we proved a non-trivial upper bound for the number of \mathcal{F} - k -facets where (X, \mathcal{F}) is any set system induced by a map. In this section we show how to improve this result for the set system $(\mathbb{R}^d, \mathcal{P}_m^d)$ for all values of m and d and for the set system $(\mathbb{R}^d, \mathcal{H}_m^d)$ for m even.

We show that the maps which induce $(\mathbb{R}^d, \mathcal{P}_m^d)$ and $(\mathbb{R}^d, \mathcal{H}_{2m}^d)$, although not neighborly, still map into convex position with a high degree of neighborliness. Since these maps come up often in many fields, the following results may be useful in other contexts.

Theorem 3.3.10. *For a finite set S of points in \mathbb{R}^d , $V_m^d(S)$ is the set of vertices of an ℓ -polytope where $\ell \leq \binom{m+d}{m} - 1$. If $V_{m/2}^d(S)$ is in general linear position and $m \geq 2$ is even, then $V_m^d(S)$ is a $\left(\binom{m/2+d}{m/2} - 1 \right)$ -neighborly ℓ -polytope.*

²This is saying that φ is generally neighborly, using the terminology of Section 4.2. However, note that φ may not be an embedding of \mathbb{R}^d . Furthermore, if $d > 2$, any map φ constructed as in Eq. (3.1) cannot be an embedding of \mathbb{R}^d .

PROOF. For $v \in S$, choose coefficients a_i, R so that $\{x \in \mathbb{R}^m : \sum_{i=1}^m (x_i - a_i)^2 \leq R\}$ is a ball with v on its boundary and with radius small enough so that no points of $S \setminus v$ are inside. By the Veronese map V_m^d , this ball corresponds to a hyperplane in $\mathbb{R}^{\binom{m+d}{m}-1}$ containing v and with all other points of S on one side. This shows that $V_m^d(v)$ is a vertex of $\text{conv}(V_m^d(S))$. For the second claim, let $T \subset S$, $|T| \leq \binom{m/2+d}{m/2} - 1$. Let $p(x) = 1$ be a degree $m/2$ polynomial passing through each point of T and no points of $S \setminus T$. To show that such a polynomial exists, recall we are assuming that $V_{m/2}^d(S)$ is in general position. Therefore, for any $|T|$ points in $V_{m/2}^d(S)$ there is a hyperplane passing through precisely those $|T|$ points. This hyperplane corresponds to a degree $m/2$ polynomial passing through each point of T and no points of $S \setminus T$. Then $(p(x) - 1)^2 = 0$ is a polynomial surface which corresponds to a hyperplane in $\mathbb{R}^{\binom{m+d}{m}-1}$ which supports $\text{conv}(T)$ as a face of $\text{conv}(S)$. \square

Theorem 3.3.11. *Assume $m \geq 2$ is even. For a finite set S of points in \mathbb{R}^d , $HV_m^d(S)$ is the set of vertices of an ℓ -polytope where $\ell \leq \binom{m+d-1}{m}$. If $HV_{m/2}^d(S)$ is in general position, meaning no hyperplane through the origin in the image space of $HV_{m/2}^d$ contains more than $\left(\binom{m/2+d-1}{m/2} - 1\right)$ points of $HV_{m/2}^d(S)$, then $HV_m^d(S)$ is a $\left(\binom{m/2+d-1}{m/2} - 1\right)$ -neighborly ℓ -polytope.*

PROOF. For $v \in S$, let $H = \{x \in \mathbb{R}^d : a \cdot x = 0\}$ be a plane through the origin which contains v and contains no other point of S . Then $(a \cdot x)^m \leq 0$ is a degree m homogeneous polynomial inequality which, by the homogeneous Veronese map HV_m^d , corresponds to a hyperplane in $\mathbb{R}^{\binom{m+d-1}{m}}$ containing v and with all other points of S on one side. This shows that $HV_m^d(v)$ is a vertex of $\text{conv}(HV_m^d(S))$. For the second claim, let $T \subset S$, $|T| \leq \left(\binom{m/2+d-1}{m/2} - 1\right)$. Let $p(x) = 0$ be a degree $m/2$ homogeneous polynomial passing through each point of T and no points of $S \setminus T$. To show that such a polynomial exists, recall we are assuming that $HV_{m/2}^d(S)$ is in general position. Therefore, for any $|T|$ points in $HV_{m/2}^d(S)$ there is a hyperplane passing through the origin and precisely those $|T|$ points. And this hyperplane corresponds to a degree $m/2$ homogeneous polynomial passing through each point of T and no points of $S \setminus T$. Then $(p(x))^2 = 0$ is a degree m homogeneous polynomial surface which corresponds to a hyperplane in $\mathbb{R}^{\binom{m+d-1}{m}}$ which supports $\text{conv}(T)$ as a face of $\text{conv}(S)$. \square

These k -neighborliness results are of interest to us because convex position is a special case of the k -set/facet problem for which we can improve the best known upper bound:

Theorem 3.3.12. *For a set S of n points in convex position in \mathbb{R}^d , $e_k(S) \leq (n/d)e_{d-1}(k, n-1)$.*

PROOF. Let $v \in S$. Choose a hyperplane H containing v and with all other points of S on one side of it. Choose another hyperplane H' parallel to H and with all points of S between H and H' . Let S' be the stereographic projection (using v as the “pole”) of $S \setminus v$ onto H' . We claim that the number of k -facets of S containing v is equal to the number of k -facets of S' (as a subset of H' , a $(d-1)$ -dimensional subspace).

Assume that $\text{conv}(v_1, \dots, v_{d-1}, v)$ is a k -facet of S . We claim that for each $s \in S$, the stereographic projection s' is on the positive side of $\text{aff}(v'_1, \dots, v'_{d-1})$ if and only if s is on the positive side of $\text{aff}(v_1, \dots, v_{d-1}, v)$. This is seen to be true by observing that $\text{aff}(s, v)$ does not intersect $\text{aff}(v_1, \dots, v_{d-1}, v)$ anywhere other than the point v . This shows that $\text{conv}(v'_1, \dots, v'_{d-1})$ is a k -facet of S' . For the converse, assume that $\text{conv}(v'_1, \dots, v'_{d-1})$ is a k -facet of S' . Then $\text{conv}(v_1, \dots, v_{d-1}, v)$ is a k -facet of S for the same reason as above.

Since S' lies in a hyperplane, it can have at most $e_{d-1}(k, n-1)$ k -facets. Performing this projection on each point of S and noticing that every k -facet is counted d times shows the desired result. \square

As far as we know, the best known bound for k -facets of n -point sets in \mathbb{R}^4 in convex position is the same as for general point sets, which is $O(n^2k^{2-2/45})$ [MSSW06]. In \mathbb{R}^3 the best known bound for general point sets is $O(nk^{2-1/2})$ [SST01a, SST01b]. This combined with Theorem 3.3.12 gives a bound of $O(n^2k^{2-1/2})$ for the number of k -facets of point sets in convex position in \mathbb{R}^4 .

An argument similar to the proof of Theorem 3.3.12 shows the following generalization (we state it without proof):

Proposition 3.3.13. *For a set S of n points in m -neighborly position in \mathbb{R}^d ,*

$$e_k(S) \leq \frac{\binom{n}{m} e_{d-m}(k, n-m)}{\binom{d}{m}}.$$

In dimensions higher than four, the best known bound for k -facets is $e_d(k, n) = O(n^{d-\epsilon_d})$ where $\epsilon_d = (4d-3)^{-d}$ [ABFK92]. Because of the fast decay of the constant ϵ_d , Proposition 3.3.13 gives

an improvement in the best known upper bound which depends on the degree of neighborliness of the point set in question.

Proposition 3.3.13 can be used to improve the bounds for \mathcal{P}_m^d - k -facets and \mathcal{H}_{2m}^d - k -facets as follows. Recall that the set system $(\mathbb{R}^d, \mathcal{P}_m^d)$ is induced by the map V_m^d and $(\mathbb{R}^d, \mathcal{H}_{2m}^d)$ is induced by the map HV_{2m}^d . Therefore Theorem 3.3.10 and Theorem 3.3.11 along with Proposition 3.3.13 give an improvement in the bound.

Generally k -neighborly embedded manifolds and varieties

To understand the limits of our argument that provides exact counting of k -facets we introduce a class of maps we call *generally neighborly embeddings* (Definition 4.2.2), which map generic point sets into neighborly position. The goal is to understand under what conditions the exact count phenomenon can occur for arbitrary set systems induced by a map. This goal leads us to ask new questions about convexity properties of embedded manifolds and algebraic varieties which are interesting beyond their connection to k -facets.

4.1. Introduction

The crucial observation that allows us to exactly count \mathcal{F} - k -facets for the conic sections and even degree homogeneous polynomials on the plane is that the maps which induce these set systems map generic point sets to neighborly point sets. In this chapter we define a *generally neighborly embedding* to be an embedding that maps generic point sets to neighborly point sets (Definition 4.2.2). The moment curve map is an example of a generally neighborly embedding of \mathbb{R}^1 . Moving up one dimension, the degree 2 Veronese map of the plane V_2^2 shows that a generally neighborly embedding of the plane exists (by Theorem 3.3.4). We conjecture (Conjecture 4.2.5) that generally neighborly embeddings of \mathbb{R}^d do not exist for $d > 2$. In order to provide support for this conjecture, in Section 4.4 we prove a closely related result about algebraic varieties. More evidence that the conjecture may be true is provided in Section 4.7.

Apart from its relation to the \mathcal{F} - k -facets, the problem of determining the existence of generally neighborly embeddings is interesting in its own right. Our definition of generally neighborly embeddings is similar to and inspired by Micha Perles' definition of *neighborly embeddings* which we discuss in Section 4.8.

Our contributions:

- **Limits of the neighborliness argument.** We study the limits of the neighborliness argument above that provides exact counting. We show that, for maps whose image is a variety, the argument only works for points on the plane. We proceed as follows: For the argument to work, one needs the map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ to map a generic set of points into a k -neighborly set of points for certain k . When φ is an embedding, we call the image $M := \varphi(\mathbb{R}^d)$ a *generally k -neighborly d -manifold* (Definition 4.3.1). We study the minimal dimension p so that M is a generally k -neighborly d -manifold and show that $p \leq 2k + d - 1$ (Theorem 4.2.4). For the same question with manifolds replaced by algebraic varieties, we show that $p = 2k + d - 1$ (Theorem 4.4.3). This line of work relates to a problem of M. Perles on k -neighborly embeddings (see Section 4.8).
- **Weakly neighborly point sets.** We leverage *weakly k -neighborly* point sets (Definition 4.5.1), a notion that is better behaved for our purposes than generally k -neighborly and Perles' k -neighborly maps (Proposition 4.5.3). In particular, we study weakly k -neighborly algebraic varieties, and resolve the question of the minimal p so that \mathbb{R}^p contains a weakly k -neighborly d -dimensional algebraic variety. We show that the minimal dimension is $2k + d - 1$ (Theorem 4.6.1).

4.2. Generally neighborly embeddings

Definition 4.2.1 (embedding). *An embedding is a map which is a homeomorphism onto its image.*

Definition 4.2.2 (generally k -neighborly embedding). *Let $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ be an embedding. For each $n \in \mathbb{N}$, let $G_n \subset \mathbb{R}^{dn}$ consist of all configurations of n points in \mathbb{R}^d which are mapped to k -neighborly sets by φ . Then φ is generally k -neighborly if G_n contains a set that is open and dense in \mathbb{R}^{dn} for all n . A generally $\lfloor \frac{p}{2} \rfloor$ -neighborly embedding is called a generally neighborly embedding.*

We choose “open and dense” in Definition 4.2.2 for concreteness and readability. For part of our discussion (in particular, Problem 4.2.3 below), it may be reasonable to substitute it by an alternative version of a generic property as discussed in Section 3.2.1.

Observe that the even degree homogeneous Veronese map of the plane, i.e. HV_m^2 , is not an embedding because it is not injective. Since the homogeneous Veronese map is one of our prime examples throughout we need to justify why we are now only talking about embeddings. The

reason is that all of the polynomial maps we have considered have the property that they are an embedding of some open subset of Euclidean space. Thus, for the purposes of this section it suffices to assume that our maps are embeddings.

The main question concerning generally k -neighborly embeddings is:

Problem 4.2.3. *What is the smallest dimension $p := p_g(k, d)$ of the image space for which a generally k -neighborly embedding $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ exists?*

Theorem 4.2.4. *There exists a generally k -neighborly embedding of \mathbb{R}^d into \mathbb{R}^{2k+d-1} and so $p_g(k, n) \leq 2k + d - 1$.*

PROOF. Consider the embedding $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^{2k+d-1}$ defined by

$$\varphi(x_1, x_2, \dots, x_d) = (x_1, x_1^2, x_1^3, \dots, x_1^{2k}, x_2, \dots, x_d).$$

For each $n \in \mathbb{N}$, let $G_n \subset \mathbb{R}^{dn}$ consist of all configurations of n points in \mathbb{R}^d such that no two points in the configuration have the same x_1 -coordinate. One can verify that G_n is open and dense in \mathbb{R}^{dn} . For any n , let $S \in G_n$ be some configuration of n points in G_n .

To show that $\varphi(S)$ is k -neighborly, let v_1, \dots, v_k be k points from S . Consider in the domain of the embedding, \mathbb{R}^d , the surface

$$(4.1) \quad \prod_{i=1}^k (x_1 - v_{i1})^2 = 0.$$

By expanding Eq. (4.1) we see that this surface corresponds via φ to a hyperplane H in \mathbb{R}^{2k+d-1} . Note that v_1, \dots, v_k satisfy Eq. (4.1) and all other points in S satisfy $\prod_{i=1}^k (x_1 - v_{i1})^2 > 0$. Using φ , we get that H is a face-defining hyperplane that makes $\varphi(v_1), \dots, \varphi(v_k)$ a face of $\text{conv}(\varphi(S))$. Therefore, $\varphi(S)$ is k -neighborly. This shows that $p_g(k, d) \leq 2k + d - 1$. \square

We believe that the bound in the above theorem is actually tight.

Conjecture 4.2.5. $p_g(k, d) = 2k + d - 1$.

Observe that for $d \geq 3$, if $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ is a generally k -neighborly embedding then, according to Conjecture 4.2.5, $p \geq 2k + 2$. This means the conjecture implies that generally neighborly embeddings of \mathbb{R}^d do not exist for $d \geq 3$.

In the context of the k -set problem, this would mean that set systems like the conic sections do not exist in dimension $d \geq 3$. More precisely, it would imply that, for $d \geq 3$, there is no set system $(\mathbb{R}^d, \mathcal{F})$ induced by an embedding (of \mathbb{R}^d) for which \mathcal{F} - k -facets can be counted by using Proposition 3.2.8.

4.3. Generally neighborly manifolds

Here we define generally k -neighborly manifolds which are, for our purposes, equivalent to generally k -neighborly embeddings. The sense in which they are equivalent is made precise below.

Definition 4.3.1. *A manifold $M \subset \mathbb{R}^p$ is generally k -neighborly if the set $G_n \subset M^n$ of configurations of n points on M which are k -neighborly contains a set that is open and dense in M^n for all n . A generally $\lfloor \frac{p}{2} \rfloor$ -neighborly manifold is called a generally neighborly manifold.*

We ask the same question for manifolds as we did for embeddings:

Problem 4.3.2. *What is the smallest dimension p of the ambient space in which a generally k -neighborly d -manifold $M \subset \mathbb{R}^p$ exists?*

Observe that an open subset of a generally k -neighborly d -manifold is still generally k -neighborly. Therefore, in the context of Problem 4.3.2 it suffices to assume that the manifold M is (globally) homeomorphic to \mathbb{R}^d , that is $M = \varphi(\mathbb{R}^d)$ for some embedding φ . This observation, along with the following proposition, shows that Problem 4.3.2 is equivalent to Problem 4.2.3.

Proposition 4.3.3. *An embedding $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ is generally k -neighborly if and only if $M := \varphi(\mathbb{R}^d)$ is a generally k -neighborly d -manifold.*

PROOF. If $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ is generally k -neighborly, then for each n the set of configurations of n points which are mapped by φ to k -neighborly point sets contains a set O that is open and dense in \mathbb{R}^{dn} . Let $N \subset M^n$ be the set of k -neighborly configurations of n -points in M^n . The set N contains $(\varphi \times \cdots \times \varphi)(O)$ which is open and dense since $\varphi \times \cdots \times \varphi$ is a homeomorphism. Therefore $M := \varphi(\mathbb{R}^d)$ is a generally k -neighborly d -manifold. The other direction is similar. \square

4.4. Generally k -neighborly algebraic varieties

Definition 4.4.1. *Let $V \subset \mathbb{R}^p$ be an irreducible real algebraic variety with a smooth real point. For each n , let $G_n \subset V_{sm}^n$ consist of all configurations of n points on V_{sm} which are k -neighborly. Then V is generally k -neighborly if G_n contains a set that is open and dense in V_{sm}^n for all n . It is generally neighborly if it is generally $\lfloor p/2 \rfloor$ -neighborly.*

We make one clarifying remark regarding the above definition. One could replace V_{sm} everywhere in the above definition with V . However, only requiring the property to hold for the smooth points strengthens our results below and does not change the proofs. Another reason for only considering the smooth points is the following. Loosely speaking, a generally k -neighborly algebraic variety V is supposed to be a variety such that every generic configuration of points on V is k -neighborly. A generic configuration of points should never contain non-smooth points, so the non-smooth points should be ignored when defining generally k -neighborly algebraic varieties.

The question we are dealing with in this section is the following.

Problem 4.4.2. *What is the smallest dimension $p := p_{g,V}(k, d)$ of the ambient space in which a generally k -neighborly d -dimensional algebraic variety $V \subset \mathbb{R}^p$ exists?*

Observe that the image of the map φ in Theorem 4.2.4 is a d -dimensional generally k -neighborly variety in \mathbb{R}^{2k+d-1} . This shows that $p_{g,V}(k, d) \leq 2k + d - 1$.

We will prove the following result which completely resolves Problem 4.4.2.

Theorem 4.4.3. *Let $V \subset \mathbb{R}^p$ be a generally k -neighborly d -dimensional algebraic variety¹. Then $p \geq 2k + d - 1$.*

Theorem 4.4.3 combined with Theorem 4.2.4 show that $p_{g,V}(k, d) = 2k + d - 1$. In order to prove Theorem 4.4.3 we will first establish a connection between generally k -neighborly varieties and *weakly k -neighborly* sets. Weak neighborliness is a more usable property that holds for all subsets of points, not just those satisfying a general position assumption. This connection is established in Section 4.5. The proof of Theorem 4.4.3 is then completed in Section 4.6.

Below we list some more examples of generally k -neighborly algebraic varieties.

¹Note that “generally k -neighborly” requires V to be irreducible and contain a smooth real point.

Example 4.4.4. *The image of the degree 2 Veronese map of the plane is a generally 2-neighborly 2-dimensional algebraic variety in \mathbb{R}^5 .*

Example 4.4.5. *The image of the map φ from the proof of Theorem 4.2.4 is a generally k -neighborly d -dimensional algebraic variety in \mathbb{R}^{2k+d-1} .*

Example 4.4.6. *The moment curve is a generally neighborly 1-dimensional algebraic variety. The same is true of any order d curve which is also an algebraic variety.*

Example 4.4.7. *Theorem 4.4.3 shows that generally neighborly d -dimensional algebraic varieties do not exist for $d \geq 3$.*

4.5. Weakly k -neighborly sets

It turns out that all generally k -neighborly algebraic varieties satisfy a weaker neighborliness property that holds for all subsets of points (not just those satisfying a general position assumption). We call this property *weakly k -neighborly* (Definition 4.5.1). In the proof of Theorem 4.4.3, we only need to use the fact that V_{sm} is weakly k -neighborly. In this section we prove some lemmas concerning weakly k -neighborly sets and the relationship between generally k -neighborly manifolds/varieties and weakly k -neighborly sets.

Definition 4.5.1. *A set $S \subseteq \mathbb{R}^p$ is weakly k -neighborly if for any set T of k points from S , there exists a closed halfspace H with boundary $\text{bd}(H)$ such that $S \subset H$ and $T \subset \text{bd}(H)$.*

We will now show that a generally k -neighborly algebraic variety or manifold is weakly k -neighborly. In order to do so, we first show that every finite subset of such a variety or manifold is weakly k -neighborly. We actually state and prove a stronger result (Lemma 4.5.2) that only uses as assumption the “dense” part of “open and dense” in the definition of generally k -neighborly. We then establish a compactness property of weakly k -neighborly sets. The property is that an arbitrary subset of \mathbb{R}^p is weakly k -neighborly if and only if every finite subset is weakly k -neighborly. These two results together establish that generally k -neighborly algebraic varieties and manifolds are weakly k -neighborly.

Lemma 4.5.2. *Let $M \subset \mathbb{R}^p$ be a manifold or the set of smooth points of an algebraic variety. If the set $N \subset M^n$ of configurations of n points on M which are k -neighborly is dense in M^n for all n , then every finite set of points on M is weakly k -neighborly.*

PROOF. Assume not, so that there exists some finite set $S \subset M$ and a set T of k points from S such that no closed halfspace contains S and contains T on its boundary. This means that $\text{aff}(T) \cap \text{relint conv}(S \setminus T) \neq \emptyset$ (from the separating hyperplane theorem [Sch14, Theorem 1.3.8]). We can pick $|S|$ small open balls in \mathbb{R}^p as follows. For each $t \in T$, let B_t be a ball centered at t and for each $s \in S \setminus T$, let A_s be a ball centered at s . The radii of the balls can be chosen small enough so that any collection of points consisting of one point from each B_t and one point from each A_s has the property that the affine hull of the points from the B_t intersects the relative interior of the convex hull of the points from the A_s . This means that any such configuration of points is not k -neighborly. Therefore, the subset of $\mathbb{R}^{p|S|}$ of configurations of points of size $|S|$ on M which are not k -neighborly contains $\prod_{t \in T} (B_t \cap M) \times \prod_{s \in S \setminus T} (A_s \cap M)$ which is an open subset of $M^{|S|}$. Therefore, the set of configurations of $|S|$ points on M which are k -neighborly is not dense in $M^{|S|}$. \square

Proposition 4.5.3 (compactness). *Let $S \subseteq \mathbb{R}^p$ be a (possibly infinite) set. Then for any $k \geq 1$ we have that S is weakly k -neighborly if and only if every finite subset of S is weakly k -neighborly.*

PROOF. Fix $k \geq 1$. The “only if” direction is clear. We will now prove the “if” direction. Let $T \subseteq S$ be a set of k points. Let $\mathcal{U} = \{U \subseteq S : U \supseteq T \text{ and } U \text{ is finite}\}$. For $U \subseteq S$ such that $U \supseteq T$, we will define $N(U) \subseteq S^{p-1}$ to be the set of unit outer normals to possible halfspaces H such that $T \subseteq \text{bd}(H)$ and $U \subseteq H$. More precisely, let

$$N(U) = \{a \in S^{p-1} : (\forall x, y \in T) a \cdot x = a \cdot y \text{ and } (\forall x \in T)(\forall y \in U) a \cdot x \geq a \cdot y\}.$$

Clearly $N(U)$ is closed. Let $\mathcal{V} \subseteq \mathcal{U}$ be any finite subfamily. Then $\cap_{U \in \mathcal{V}} N(U) = N(\cup_{U \in \mathcal{V}} U) \neq \emptyset$ by assumption. We have established that $\{N(U)\}_{U \in \mathcal{U}}$ is a family of closed sets with the finite intersection property in compact space S^{p-1} . This implies $\cap_{U \in \mathcal{U}} N(U) \neq \emptyset$. We also have $N(S) = N(\cup_{U \in \mathcal{U}} U) = \cap_{U \in \mathcal{U}} N(U) \neq \emptyset$. That is, there is a halfspace H such that $T \subseteq \text{bd}(H)$ and $S \subseteq H$. As $T \subseteq S$ was arbitrary, this completes the proof. \square

We require two more lemmas concerning intersections and weak separation of convex sets in \mathbb{R}^p . (See Section 2.1 for the definitions of weak and proper separation of sets in \mathbb{R}^d .)

Lemma 4.5.4 (Radon-type theorem). *Let P be a set of $p+2$ points in \mathbb{R}^p in general linear position. Then there is a partition Q, R of P into two non-empty sets so that $\text{relint conv } Q \cap \text{relint conv } R \neq \emptyset$.*

PROOF. Let $P = \{q_1, \dots, q_{p+2}\}$. P is affinely dependent and therefore there exist $\lambda_1, \dots, \lambda_{p+2}$ such that $\sum_{i=1}^{p+2} \lambda_i q_i = 0$, $\sum_{i=1}^{p+2} \lambda_i = 0$ and at least one λ_i is non-zero. Because of the general position assumption, all λ_i are non-zero. Let $I = \{i : \lambda_i > 0\}$, $J = \{i : \lambda_i < 0\}$. Both I and J are non-empty. By dividing λ_i s by $\sum_{i \in I} \lambda_i$ we can assume without loss of generality that $\sum_{i \in I} \lambda_i = -\sum_{i \in J} \lambda_i = 1$. Let $Q = \{q_i : i \in I\}$, $R = \{q_i : i \in J\}$. Let $q = \sum_{i \in I} \lambda_i q_i \in \text{relint conv } Q$, $r = -\sum_{i \in J} \lambda_i q_i \in \text{relint conv } R$. We have $q = r$, which completes the proof. \square

Lemma 4.5.5. *Let $Q, R \subseteq \mathbb{R}^p$ be disjoint sets. Suppose $\text{aff}(Q \cup R) = \mathbb{R}^p$ and $\text{relint conv } Q \cap \text{relint conv } R \neq \emptyset$. Then Q, R cannot be weakly separated.*

PROOF. Assume Q, R can be weakly separated. If the separation is not proper then $\text{aff}(Q \cup R) \neq \mathbb{R}^p$. If the separation is proper, then by the separating hyperplane theorem [Sch14, Theorem 1.3.8], $\text{relint conv } Q \cap \text{relint conv } R = \emptyset$. \square

4.6. Weakly k -neighborly varieties

In the previous section we established the connection between generally k -neighborly varieties/manifolds and weakly k -neighborly sets. In this section, we use this connection to prove Theorem 4.4.3. Given an (real) algebraic variety V of dimension d such that V_{sm} is weakly k -neighborly, we prove a sharp lower bound on the dimension of the ambient space.

Theorem 4.6.1. *Assume $V \subset \mathbb{R}^p$ is a non-degenerate d -dimensional irreducible real algebraic variety with a smooth real point. If $V \setminus U$ is weakly k -neighborly for some proper closed subvariety U , then $p \geq 2k + d - 1$.*

Before proving Theorem 4.6.1 we prove the special case of algebraic curves and then generalize to higher dimensional varieties.

Lemma 4.6.2. *Assume $C \subset \mathbb{R}^p$ is a non-degenerate irreducible real algebraic curve with a smooth real point. If $C \setminus U$ is weakly k -neighborly for some proper closed subvariety U , then $p \geq 2k$.*

PROOF. First we will show that one can find arbitrarily large point sets in general linear position on $C \setminus U$. In order to accomplish this, first observe that $C \setminus U$ is non-degenerate. Indeed, if it were the case that $C \setminus U$ is contained in a hyperplane H , then we would have $C = (C \cap H) \cup U$ which is impossible since C is irreducible.

Now assume that S is a set of j points in general linear position on $C \setminus U$. If it were not possible to find another point s such that $S \cup \{s\}$ is in general linear position, it would have to be the case that $C \setminus U$ is contained in the union of all hyperplanes spanned by p points from the set S . Let \mathcal{H} be the collection of all hyperplanes spanned by p points in S . Note that \mathcal{H} is finite. We have

$$C = \left(\bigcup_{H \in \mathcal{H}} ((C \setminus U) \cap H) \right) \cup U.$$

Since $C \setminus U$ is non-degenerate the above formula would be a representation of C as the union of proper subvarieties. This is impossible since C is irreducible. Therefore we know that $C \setminus U$ is not contained in the union of all hyperplanes spanned by points in S and so we can always find s so that $S \cup \{s\}$ is in general linear position. It follows that we can find arbitrarily large point sets in general linear position on $C \setminus U$.

Let P be a set of $p + 2$ points in general linear position on $C \setminus U$. By Lemmas 4.5.4 and 4.5.5, there is a partition of P into non-empty sets Q and R so that Q and R cannot be weakly separated.

However, because $C \setminus U$ is weakly k -neighborly, we know that for any set T of k points on $C \setminus U$ there exists a closed halfspace which contains $C \setminus U$ and contains T in its boundary. In other words, any such T can be weakly separated from $C \setminus U$. Therefore, it must be that $k < \min(|Q|, |R|)$, i.e. $k \leq \min(|Q|, |R|) - 1$. Now since $\min(|Q|, |R|) \leq \lfloor \frac{p+2}{2} \rfloor$, we have that $k \leq \lfloor \frac{p+2}{2} \rfloor - 1$ and so $p \geq 2k$. \square

The idea of the proof for higher dimensional varieties is to take successive hyperplane sections in order to reduce to the case of curves. So we first need to establish the following lemma concerning hyperplane sections of varieties.

Lemma 4.6.3. *Assume that $V \subset \mathbb{R}^p$ is a non-degenerate d -dimensional ($d \geq 2$) irreducible variety with a smooth real point and that U is some proper closed subvariety of V . Then for any given open ball B in $\text{conv}(V \setminus U)$ there exists a hyperplane H such that $H \cap B \neq \emptyset$ and $V \cap H$ is a non-degenerate, irreducible, $(d - 1)$ -dimensional variety with a smooth real point which is not contained in U .*

PROOF. We identify the set of hyperplanes in \mathbb{R}^p with a proper subset of $\mathbb{P}^p(\mathbb{R})$. This identification works as follows. We identify a hyperplane $a_0 + a_1x_1 + \cdots + a_px_p = 0$ in \mathbb{R}^p with the point with homogeneous coordinates (a_0, a_1, \dots, a_p) in $\mathbb{P}^p(\mathbb{R})$. Therefore, the set of hyperplanes in \mathbb{R}^p is identified with $\mathbb{P}^p(\mathbb{R}) \setminus \{(a_0, a_1, a_2, \dots, a_p) : a_1 = a_2 = \cdots = a_p = 0\}$.

Recall from Section 2.5 that around any smooth real point of V there is a neighborhood which is a smooth d -dimensional submanifold of \mathbb{R}^p . Let $T \subset \mathbb{P}^p(\mathbb{R})$ be the set of hyperplanes H for which there exists a smooth point of $V \setminus U$ and a neighborhood of that smooth point which has nonempty transversal intersection with H . We know that T is open and that for any $H \in T$, $H \cap V$ contains an open subset which is a smooth $(d - 1)$ -dimensional submanifold of \mathbb{R}^p [GP10, Sections 1.5 and 1.6]. Let H be any hyperplane in T . Clearly the dimension of $H \cap V$ is at most $d - 1$. We claim that $H \cap V$ is a variety of dimension precisely $d - 1$ and that it contains a smooth real point. Recall from Section 2.5 that if a variety has real dimension d then it has dimension at least d . Therefore, the dimension of $H \cap V$ is $d - 1$. Now we show that $H \cap V$ contains a smooth real point. Indeed assume not, so that $H \cap V(\mathbb{C})$ contains no smooth real points.² This means that $H \cap V$ is contained in the set of singular points of the $(d - 1)$ -dimensional variety $H \cap V(\mathbb{C})$. By [BCR98, Proposition 3.3.14], the set of singular points of $H \cap V(\mathbb{C})$ is a variety of dimension at most $d - 2$. Since $H \cap V$ has real dimension $d - 1$ and is contained in a variety of dimension at most $d - 2$ this is a contradiction. So $H \cap V$ contains smooth real points.

Let I be the subset of T consisting of hyperplanes H such that $H \cap V$ is irreducible and non-degenerate. We will show that I is open and dense in the standard topology in T . Let $\overline{V(\mathbb{C})} \subset \mathbb{P}^p(\mathbb{C})$ be the projective closure of $V(\mathbb{C})$. This means that $V(\mathbb{C}) = \overline{V(\mathbb{C})} \cap \{x_0 \neq 0\}$. Because V contains a smooth real point, V is Zariski dense in $V(\mathbb{C})$, that is, $V(\mathbb{C})$ is the smallest complex variety

²By a minor abuse of notation, $H \cap V(\mathbb{C})$ is the complex variety defined by the polynomials defining V along with the polynomial defining H . Similarly, $H \cap V$ is the real variety defined by the polynomials defining V along with the polynomial defining H . So $H \cap V$ is the real part of $H \cap V(\mathbb{C})$.

containing V , see [BCR98, Section 2.8]. Therefore, since V is irreducible, by [Whi57, Lemma 7], $V(\mathbb{C})$ is irreducible. Because $V(\mathbb{C})$ is irreducible, it is then a standard fact that $\overline{V(\mathbb{C})}$ is irreducible. By projective duality, we can identify the set of all projective hyperplanes with real coefficients in $\mathbb{P}^p(\mathbb{C})$ with $\mathbb{P}^p(\mathbb{R})$. Let $I' \subset \mathbb{P}^p(\mathbb{R})$ consist of all projective hyperplanes with real coefficients that have irreducible and non-degenerate intersection with $\overline{V(\mathbb{C})}$. By [Har95, Theorem 18.10], I' is Zariski open and dense.³

We will now show that the fact that I' is Zariski open and dense implies that I is also Zariski open and dense in T . Given a hyperplane H defined by $a_0 + a_1x_1 + \cdots + a_px_p = 0$ in \mathbb{R}^p , there is a corresponding projective hyperplane \overline{H} defined by $a_0x_0 + a_1x_1 + \cdots + a_px_p = 0$ which is the homogenization of H . Let H be a hyperplane in T and assume that the homogenization \overline{H} is in I' . We claim that this implies that $H \in I$. To establish this claim, we need to verify that $H \cap V$ is irreducible and nondegenerate. Observe that $H \cap V(\mathbb{C}) = \overline{H} \cap \overline{V(\mathbb{C})} \cap \{x_0 \neq 0\}$ i.e., $H \cap V(\mathbb{C})$ is an open subset of $\overline{H} \cap \overline{V(\mathbb{C})}$. It is a standard fact that a nonempty open subset of an irreducible space is irreducible and dense. Therefore $H \cap V(\mathbb{C})$ is irreducible. Now, since $H \cap V(\mathbb{C})$ is Zariski dense in $\overline{H} \cap \overline{V(\mathbb{C})}$, if $H \cap V(\mathbb{C})$ were contained in a hyperplane, $\overline{H} \cap \overline{V(\mathbb{C})}$ would be as well. So $H \cap V(\mathbb{C})$ is nondegenerate. We have shown that $H \cap V(\mathbb{C})$ is non-degenerate and irreducible. We claim that because the section $H \cap V(\mathbb{C})$ contains smooth real points, then the non-degeneracy and irreducibility of $H \cap V(\mathbb{C})$ implies non-degeneracy and irreducibility of $H \cap V$. This relies on the fact mentioned above that if an irreducible variety W contains a smooth real point, then the set of real points is Zariski dense in $W(\mathbb{C})$. This means that $H \cap V$ is Zariski dense in $H \cap V(\mathbb{C})$. Now by [Whi57, Lemma 7], irreducibility of $H \cap V(\mathbb{C})$ implies that $H \cap V$ is irreducible. Finally, observe that because $H \cap V$ is Zariski dense in $H \cap V(\mathbb{C})$, if $H \cap V$ is degenerate, $H \cap V(\mathbb{C})$ must be as well. Thus $H \cap V$ is non-degenerate.

Recall that T is open (in the standard metric topology) and that I' is Zariski open and dense in $\mathbb{P}^p(\mathbb{R})$ which implies that I' is open and dense in the standard topology. We established that I contains $I' \cap T$. Therefore, we have shown that I is open and dense in the standard topology in T .

³Theorem 18.10 in [Har95] says that the set of hyperplanes which intersect $\overline{V(\mathbb{C})}$ in a non-degenerate irreducible variety is the complement of a proper subvariety of $\mathbb{P}^p(\mathbb{C})$. The intersection of a proper subvariety of $\mathbb{P}^p(\mathbb{C})$ with $\mathbb{P}^p(\mathbb{R})$ is a proper subvariety of $\mathbb{P}^p(\mathbb{R})$. So I' is Zariski open dense in $\mathbb{P}^p(\mathbb{R})$.

Let $O \subset \mathbb{P}^p(\mathbb{R})$ be the set of hyperplanes which intersect B . The set O is open (in the standard topology).

We claim that all of this means that $I \cap T \cap O$ is non-empty. To establish this, we need to verify that $T \cap O$ is non-empty. To find $H \in T \cap O$, let v be a smooth point of $V \setminus U$ and let F be a $(p - 2)$ -flat having non-empty transversal intersection with a neighborhood of v . Such a flat exists because V has real dimension at least 2. For any point p in B , the hyperplane $\text{aff}(F \cup p)$ is in $T \cap O$.

Therefore, $T \cap O$ is open and non-empty. Since I is open and dense in T , it follows that $I \cap T \cap O \neq \emptyset$. We claim that any $H \in I \cap T \cap O$ completes the proof. To show this, it remains to show that for any $H \in I \cap T \cap O$, $H \cap V$ has a smooth real point which is not in U . We already know that $H \cap V$ has smooth real points and that $H \cap V$ is not contained in U . So assume for a contradiction that all the smooth real points are contained in U . Then letting S denote the singular points of $H \cap V$, we have that $H \cap V = (H \cap V \cap U) \cup S$ is the decomposition of $H \cap V$ as the union of two proper closed subvarieties, a contradiction to irreducibility of $H \cap V$. \square

The lemma established above allows us to generalize Lemma 4.6.2 to higher dimensional varieties.

PROOF OF THEOREM 4.6.1. First we show that by making repeated applications of Lemma 4.6.3, we can inductively construct a $(p - d + 1)$ -flat L that intersects $\text{int conv}(V \setminus U)$ and such that $L \cap V$ is a non-degenerate irreducible algebraic curve with a smooth real point which is not contained in U . Say we have some flat F that intersects $\text{int conv}(V \setminus U)$ and such that $F \cap V$ is a non-degenerate irreducible d' -dimensional ($d' \geq 2$) variety with a smooth real point not contained in U . Since $F \cap V$ is not contained in U , $F \cap U$ is a proper subvariety of $F \cap V$. Let B be an open ball in F that is contained in $\text{int conv}(V \setminus U)$. By Lemma 4.6.3, there exists a hyperplane H in F that intersects B and such that $H \cap V$ is a non-degenerate irreducible $(d' - 1)$ -dimensional variety with a smooth real point not contained in $F \cap U$ and hence not contained in U . Since B is contained in $\text{int conv}(V \setminus U)$, H intersects $\text{int conv}(V \setminus U)$. We can repeat this process until we obtain a $(p - d + 1)$ -flat L that intersects $\text{int conv}(V \setminus U)$ and such that $C := L \cap V$ is a non-degenerate irreducible algebraic curve with a smooth real point not contained in U .

Now we will show that $C \setminus U$ is weakly k -neighborly in L , meaning⁴ that $C \setminus U$ is weakly k -neighborly as a subset of its affine hull L .

Because $V \setminus U$ is weakly k -neighborly, we know that for any set T of k points on $C \setminus U$ there exists a closed halfspace H (in \mathbb{R}^p) which contains $C \setminus U$ and contains T in $\text{bd}(H)$. In other words, any such T can be weakly separated from $C \setminus U$. Notice that we are talking about weak separation in \mathbb{R}^p , while we are really interested in weak separation in L . We claim that our assumption that L intersects the interior of the convex hull of $V \setminus U$ allows us to pass from weakly separating hyperplanes in \mathbb{R}^p to weakly separating hyperplanes in L . Indeed, the fact that L intersects $\text{int conv}(V \setminus U)$ means that any closed halfspace H satisfying $V \setminus U \subset H$ cannot contain L in its boundary. Therefore $\text{bd}(H) \cap L$ is a proper hyperplane in L . To summarize, given a set T of k points on $C \setminus U$, by the assumption that any finite set on $V \setminus U$ is weakly k -neighborly, there exists a hyperplane $\text{bd}(H)$ that weakly separates T from $V \setminus U$. This and the way we chose L allows us to conclude that $\text{bd}(H) \cap L$ is a hyperplane in L that weakly separates T from $C \setminus U$. Therefore, $C \setminus U$ is weakly k -neighborly in L . Since C is not contained in U , $C \cap U$ is a proper subvariety. By Lemma 4.6.2, $p - d + 1 \geq 2k$. \square

We can now prove the lower bound for generally k -neighborly varieties.

PROOF OF THEOREM 4.4.3. We can without loss of generality assume that V is non-degenerate since otherwise we could consider V in $\text{aff}(V)$. By Lemma 4.5.2, every finite set of points on V_{sm} is weakly k -neighborly. Therefore by Proposition 4.5.3, V_{sm} is weakly k -neighborly. Since V_{sing} is a proper closed subvariety (see [BCR98, Proposition 3.3.14]), by Theorem 4.6.1, $p \geq 2k + d - 1$. \square

4.7. Additional evidence

In this section, we give some additional comments on the validity of Conjecture 4.2.5. Although we could not resolve the conjecture, the result on algebraic varieties is evidence that it is likely true. In the following we provide more evidence for the conjecture by showing that any manifold violating the conjecture would have to have a property that appears fairly restrictive to us.

⁴Note that a subset of \mathbb{R}^p whose affine hull is a proper subset of \mathbb{R}^p is automatically weakly k -neighborly according to Definition 4.5.1, so the requirement here is that halfspace H in that definition is a halfspace of L .

Proposition 4.7.1. *If a set $M \subset \mathbb{R}^p$ is weakly k -neighborly then for any set S of $2k$ points in general linear position in M , $\text{aff}(S) \cap M$ is contained in the union of all hyperplanes supporting facets of the simplex $\text{conv } S$.*

PROOF. Assume not, so that S is a set of $2k$ points in M such that $\text{aff}(S) \cap M$ is not contained in the union of all hyperplanes supporting facets of $\text{conv}(S)$. Then identifying $\text{aff}(S)$ with \mathbb{R}^{2k-1} , we can find a set S' of $2k + 1$ points in general position in \mathbb{R}^{2k-1} which are weakly k -neighborly. However, by Lemmas 4.5.4 and 4.5.5, there is a partition Q, R of S' into two non-empty sets so that Q, R cannot be weakly separated. Since $\min(Q, R) \leq \lfloor \frac{2k+1}{2} \rfloor = k$, this is a contradiction to weakly k -neighborliness of S' . \square

Proposition 4.7.1 is inspired by the following illustrative example: the possibility of a non-degenerate 2-neighborly curve in \mathbb{R}^3 . By non-degeneracy, pick 4 affinely independent points on the curve. Then by Proposition 4.7.1 the curve would have to be contained in the union of the 4 hyperplanes defined by any 3 of those points.

Assume $M \subset \mathbb{R}^p$ is a weakly k -neighborly embedded d -manifold and S any set of $2k$ points in general position on M . Notice that if $p < 2k + d - 1$, we would expect $\text{aff}(S)$ to intersect M in a manifold of dimension 1 or greater. However, the previous Proposition implies that $M \cap \text{aff } S$ is contained in a finite number of hyperplanes which is not true of most embedded 1-manifolds.

One approach to proving Conjecture 4.2.5 appears to be showing that, in fact, there is no non-degenerate d -manifold in \mathbb{R}^{2k+d-2} satisfying the conclusion of Proposition 4.7.1.

4.8. Neighborly embeddings

Our definition of generally k -neighborly embeddings is similar to the concept of *k -neighborly embeddings* introduced by Perles in 1982 and studied by Kalai and Wigderson [KW08].

An embedding of a d -dimensional manifold M into \mathbb{R}^p is *k -neighborly* if for every k points on the embedding of M there is a hyperplane H that contains the k points and such that all remaining points of the embedded manifold are on the (strictly) positive side of H .

Requiring that an embedding be k -neighborly is clearly stronger than requiring that it be generally k -neighborly. That is, if an embedding is k -neighborly then it is generally k -neighborly. However, the reverse implication is certainly not true. For example, the degree two Veronese

embedding $V_2^2(x, y) = (x, y, x^2, xy, y^2)$ is only a 1-neighborly embedding while it is a generally 2-neighborly embedding.

In 1982 Perles posed the following problem concerning neighborly embeddings.

Problem 4.8.1. What is the smallest dimension $p(k, d)$ of the ambient space in which a k -neighborly d -dimensional manifold exists?

As in the case of generally neighborly embeddings, for the purposes of this question, it suffices to assume that $M = \mathbb{R}^d$. Kalai and Wigderson proved

THEOREM 4.8.2 ([KW08]). $k(d + 1) \leq p(k, d) \leq 2k(k - 1)d$.

Improving the bounds in Theorem 4.8.2 appears to be difficult compared to the case of generally k -neighborly embeddings where we were able to conjecture a precise formula for $p_g(k, d)$.

Comparing the two definitions, k -neighborly embeddings appear to be the more natural and fundamental class of embeddings to investigate. However, there may be some applications for which the notion of generally k -neighborly embeddings is more appropriate. For example, the authors of [KW08] were interested in neighborly embeddings in part because they may lead us to important examples of k -neighborly polytopes. In particular, by picking points on the embedded manifold, one may produce k -neighborly nonsimplicial polytopes perhaps with other interesting properties. Since both types of embeddings produce k -neighborly polytopes, our version may be more useful in this context as it is less restrictive.

Improved bounds for the expected number of k -sets

For a probability distribution P on \mathbb{R}^d , we study $E_P(k, n)$, the expected number of k -facets of a sample of n random points from P . When P is a distribution on \mathbb{R}^2 such that the measure of every line is 0, we show that $E_P(k, n) = O(n^{5/4})$ when $k = \lfloor cn \rfloor$ for any fixed $c \in (0, 1)$. Our argument is based on a technique by Bárány and Steiger. We study how it may be possible to improve this bound using the continuous version of the polynomial partitioning theorem. This motivates a question concerning the points of intersection of an algebraic curve and the k -edge graph of a set of points.

5.1. Expected number of k -edges

Recall that $e_2(k, n)$ denotes the maximum number of k -edges of any set of n points in the plane. It is widely believed that the true value of $e_2(k, n)$ is closer to the best known lower bound than to the best known upper bound. Indeed, Erdős et al. conjectured in [ELSS73] that $e_2(k, n) = O(n^{1+\epsilon})$ for any $\epsilon > 0$. Some support for this conjecture is provided by the results one can obtain for the probabilistic version of the k -facet problem.

Bárány and Steiger initiated the study of the probabilistic version of the k -facet problem [BS94]. The problem was also studied in [Cla04]. Given a probability distribution P on \mathbb{R}^d , what is the expected number $E_P(k, n)$ of k -facets of X , a sample of n independent random points from P ? Recall that the k -facet problem is only defined for point sets in general position. For this reason, in all of our results concerning $E_P(k, n)$, we restrict our attention to distributions P such that the measure of every hyperplane is 0. This is the minimal assumption on distributions P which guarantees that a sample of points from P is in general position with probability 1.

We refer to the original k -facet problem as the *deterministic* version. By the *probabilistic* version of the k -facet problem, we mean the question of the value of $E_P(k, n)$.

Part of the reason for the interest in the probabilistic version of the k -facet problem is that, as noted in [Wag08, Section 4.2], there is some reason to believe that the probabilistic version may be more or less the same as the original: Circa publication of Bárány and Steiger's paper [BS94], the best known lower bound for the deterministic k -facet problem in \mathbb{R}^2 was $e_d(\frac{n-2}{2}, n) = \Omega(n \log n)$ [ELSS73]. Using the construction in [ELSS73], Bárány and Steiger construct a probability distribution P with $E_P(\frac{n-2}{2}, n) = \Omega(n \log n)$. The $\Omega(n \log n)$ lower bound for the deterministic k -facet problem was improved to $e_d(\frac{n-2}{2}, n) = ne^{\Omega(\sqrt{\log n})}$ in [Tót01]. As noted in [Wag08], it is possible to use the construction in [Tót01] to construct a distribution P' with $E_{P'}(\frac{n-2}{2}, n) = ne^{\Omega(\sqrt{\log n})}$. See [Wag08] for some more details.

Upper bounds for the probabilistic k -facet problem which improve the upper bounds for the deterministic version have only been established in special cases. Bárány and Steiger obtained tight upper bounds for several families of distributions using an integral formula for $E_P(k, n)$. We also use the formula in the proofs of our main results so we describe it here. Let X_1, \dots, X_d be d random points drawn from P . The assumption that the measure of every hyperplane is zero means that $\text{aff}(X_1, \dots, X_d)$ is a hyperplane with probability 1. We use $\text{aff}(X_1, \dots, X_d)^+$ to denote the open half space on the positive side of $\text{aff}(X_1, \dots, X_d)$ and $\text{aff}(X_1, \dots, X_d)^-$ to denote the open half space on the negative side of $\text{aff}(X_1, \dots, X_d)$. Define

$$G(t) = \mathbb{P}(P(\text{aff}(X_1, \dots, X_d)^+) \leq t).$$

Then given a sample X of n random points from P , the expected number of k -facets of X is

$$\begin{aligned} E_P(k, n) &= \sum_{F \in \binom{X}{d}} \mathbb{P}(\text{aff}(F)^+ \text{ or } \text{aff}(F)^- \text{ contains exactly } k \text{ points of } X) \\ (5.1) \quad &= 2 \binom{n}{d} \binom{n-d}{k} \int_0^1 t^k (1-t)^{n-d-k} dG(t). \end{aligned}$$

Bárány and Steiger used Eq. (5.1) to show that $E_P(k, n) = O(n^{d-1})$ if P is spherically symmetric. Also, if P is the uniform distribution on a convex body in \mathbb{R}^2 , they show that $E_P(k, n) = O(n)$.

Finally, we remark that Eq. (5.1) can be used to immediately obtain an upper bound for the expected number of $\frac{n-d}{2}$ -facets of a sample of n points from *any*¹ probability distribution on \mathbb{R}^d .

¹Recall we assume that hyperplanes have measure 0.

The bound is much weaker than the bounds obtained by Bárány and Steiger for special cases of the distribution, but it is still non-trivial.

Theorem 5.1.1. *If P is a Borel probability distribution on \mathbb{R}^d such that the measure of any hyperplane is zero, then $E_P(\frac{n-d}{2}, n) = O(n^{d-1/2})$.*

PROOF. From Eq. (5.1) and Stirling's approximation,

$$\begin{aligned}
 E_P\left(\frac{n-d}{2}, n\right) &= 2 \binom{n}{d} \binom{n-d}{\frac{n-d}{2}} \int_0^1 t^k (1-t)^{n-d-k} dG(t) \\
 (5.2) \qquad \qquad \qquad &\leq 2 \binom{n}{d} \binom{n-d}{\frac{n-d}{2}} \frac{1}{2^{n-d}} \\
 &= O(n^{d-1/2}).
 \end{aligned}$$

□

In Section 5.2 we use the weak bound in the proof of Theorem 5.1.1 combined with a partition of the plane to show an improved bound for the expected number of k -edges:

Theorem 5.1.2. *Let P be a Borel probability distribution on \mathbb{R}^2 such that the measure of every line is zero. Then for any $1 \leq k \leq n-1$, $E_P(k, n) = O\left(\frac{n^{7/4}}{k^{1/4}(n-k)^{1/4}}\right)$ (where the constants in big- O are universal).*

The proof of Theorem 5.1.2 is in Section 5.2.2. In Section 5.2.1, we review the definition of the k -edge graph (Definition 5.2.1) and state some results needed for the proof of Theorem 5.1.2. In particular, we explain how the k -edge graph can be decomposed into *convex chains*. The proof of Theorem 5.1.2 also uses a divide and conquer approach where the plane is partitioned into cells by vertical lines.

Section 5.3 outlines how it may be possible to improve the bound in Theorem 5.1.2 by partitioning the plane with algebraic curves rather than vertical lines. The existence of algebraic curves which partition the plane in a useful way is a consequence of the continuous version of the polynomial partitioning theorem of [Gut16b] (Theorem 2.3.4). Whether or not using algebraic curves rather than lines leads to an improvement in the bound depends on a question we leave open (Question 5.3.1) which asks for a bound on the maximum number of times that an algebraic

curve of degree r can intersect the k -edge graph of a set of n points. We show that this quantity is $O(nr^2)$ but, as far as the authors know, it may be possible to improve this bound, see Section 7.4.

5.2. Bounding the expected number of k -edges

In this section we review some necessary facts about k -edges and prove Theorem 5.1.2.

5.2.1. Convex/concave chains. Here we recall the “convex chains” technique of [AACS98] which was used in [Dey97] to establish the $O(n^{4/3})$ bound for planar k -edges. First we need to define the k -edge graph.

Let S be a set of n points in general position in the plane and choose some (x, y) coordinate system. With this choice, we assume without loss of generality that no line spanned by two points in S is vertical. Let E_k be the set of line segments connecting two points $x, y \in S$ such that there are exactly k points from S in the halfplane below $\text{aff}(x, y)$. Therefore, E_k is a subset of the set of all k -edges of S . Throughout, we assume without loss of generality that E_k contains at least half the total number of k -edges. Indeed, we could repeat the analysis after rotating the plane 180 degrees. The line segments E_k define the k -edge graph:

Definition 5.2.1. *Let S be a set of n points in general position in the plane and assume that no line spanned by two points in S is vertical. For any $0 \leq k \leq n - 2$, the (geometric) graph $G_k = (S, E_k)$ is called the k -edge graph of S .*

The convex chains technique decomposes the k -edge graph G_k of a point set S into the union of a bounded number of *convex chains*. Each convex chain is the graph of a convex piece-wise linear function defined on some interval of the x -axis. Each chain is formed by some subset of the k -edges in $G_k = (S, E_k)$. A simpler version of the proof of the $O(n^{4/3})$ bound was established in [HP11] by observing that the k -edge graph can simultaneously be decomposed into the union of *concave chains*. We use the known fact [Lov71] that the convex/concave chain decompositions imply that a vertical line can intersect the k -edge graph at most $\min(k - 1, n - k - 1)$ times.

Lemma 5.2.2 (Convex/concave chains [HP11, Lemma 9.10]). *Let $G_k = (S, E_k)$ be the k -edge graph of a set S of n points in the plane. The graph G_k can be decomposed into the union of $k - 1$*

(piece-wise linear) convex chains. Similarly, the graph can be decomposed into the union of $n - k + 1$ (piece-wise linear) concave chains.

These decompositions can be used to show that the total number of crossings in G_k is $O(n^2)$, see [Dey97] or [HP11].

5.2.2. Proof of Theorem 5.1.2. We are now ready to prove our bound on the expected number of k -edges. The idea of the proof is to use vertical lines to divide the plane into a number of regions of equal probability. We then bound separately the expected number of k -edges that intersect one of the vertical lines and the expected number of k -edges that do not intersect any of the lines. We remark that a partition of the plane using vertical lines was also used by Lovász in [Lov71] to establish the $O(n^{3/2})$ bound for the deterministic k -set problem.

PROOF OF THEOREM 5.1.2. For any fixed n , let $m := m(n)$ be an integer whose value will be chosen later and let L be the set consisting of m vertical lines which partition $\mathbb{R}^2 \setminus (\bigcup_{\ell \in L} \ell)$ into $m + 1$ open cells such that the measure according to P of each cell is equal to $1/(m + 1)$. Such a set of lines exists because the measure (according to P) of every line is 0. In order to show how to construct this set of lines, it suffices to show that given any finite measure μ on \mathbb{R}^2 such that the measure of every line is 0, there exists a vertical line which divides the measure into two equal parts. Indeed, the function f defined by $f(t) = \mu((x, y) \in \mathbb{R}^2 : x \leq t)$ is a continuous function of t . The fact that f is continuous is a consequence of the assumption that the measure of every vertical line is 0. Furthermore, as $t \rightarrow -\infty$, $f(t) \rightarrow 0$ and as $t \rightarrow \infty$, $f(t) \rightarrow \mu(\mathbb{R}^2)$. Therefore, the claim follows from the intermediate value theorem.

Let $X = \{X_1, \dots, X_n\}$ be a sample of n iid points from P . Observe that the probability that two points in X span a vertical line is 0 so we do not need to consider this case. Also, since the measure of every line is 0, X is in general position with probability 1 so we can also ignore the case when X is not in general position. Therefore, we can analyze the k -edges of X using the k -edge graph G_k of X .

We bound the expected number of k -edges in G_k by considering two different types of k -edges separately. First we bound the expected number of k -edges formed by two points in different cells of the partition. Then we bound the expected number of k -edges formed by two points in the same

cell of the partition. That is, the expected number of k -edges in G_k is equal to

$$(5.3) \quad \begin{aligned} & \mathbb{E}(\text{number of } k\text{-edges in } G_k \text{ formed by two points in different cells}) \\ & + \mathbb{E}(\text{number of } k\text{-edges in } G_k \text{ formed by two points in the same cell}). \end{aligned}$$

If $\text{conv}(X_1, X_2)$ is a k -edge formed by two points X_1, X_2 in different cells, then $\text{conv}(X_1, X_2)$ intersects at least one line in L . So to bound the expected number of k -edges in G_k formed by two points in different cells, it suffices to bound the expected number of k -edges in G_k that intersect a line in L . By Lemma 5.2.2, each line in L intersects at most $\min(k-1, n-k-1)$ k -edges in G_k . Therefore, the first term in Eq. (5.3) is at most $m \cdot \min(k-1, n-k-1)$

Now we bound the second term. Recall that the measure according to P of each cell is equal to $1/(m+1)$. Therefore, for any fixed $i \neq j$, the probability that X_i and X_j are in the same cell is $1/(m+1) \leq 1/m$.

We can bound the second term in Eq. (5.3) by

$$(5.4) \quad \begin{aligned} & \mathbb{E}(\text{number of } k\text{-edges in } G_k \text{ formed by two points in the same cell}) \\ & \leq \sum_{(X_i, X_j) \in \binom{X}{2}} \mathbb{P}((X_i, X_j) \text{ is a } k\text{-edge AND } X_i, X_j \text{ are in the same cell}) \\ & = \sum_{(X_i, X_j) \in \binom{X}{2}} \mathbb{P}((X_i, X_j) \text{ is a } k\text{-edge} \mid X_i, X_j \text{ are in the same cell}) \\ & \quad \cdot \mathbb{P}(X_i, X_j \text{ are in the same cell}) \\ & = \binom{n}{2} \cdot \mathbb{P}((X_1, X_2) \text{ is a } k\text{-edge} \mid X_1, X_2 \text{ are in the same cell}) \\ & \quad \cdot \mathbb{P}(X_1, X_2 \text{ are in the same cell}) \\ & \leq \frac{n^2}{m} \cdot \mathbb{P}((X_1, X_2) \text{ is a } k\text{-edge} \mid X_1, X_2 \text{ are in the same cell}). \end{aligned}$$

Set $T := P(\text{aff}(X_1, X_2)^+)$ and $G(t) = \mathbb{P}(T \leq t \mid X_1, X_2 \text{ are in the same cell})$. We then have that

$$\begin{aligned}
& \mathbb{P}((X_1, X_2) \text{ is a } k\text{-edge} \mid X_1, X_2 \text{ are in the same cell}) \\
&= \mathbb{E}\left(\mathbb{P}((X_1, X_2) \text{ is a } k\text{-edge} \mid X_1, X_2, (X_1, X_2 \text{ are in the same cell}))\right. \\
&\quad \left. \mid X_1, X_2 \text{ are in the same cell}\right) \\
(5.5) \quad &= 2 \binom{n-2}{k} \mathbb{E}(T^k (1-T)^{n-2-k} \mid X_1, X_2 \text{ are in the same cell}) \\
&= 2 \binom{n-2}{k} \int_0^1 t^k (1-t)^{n-2-k} dG(t) \\
&\leq 2 \binom{n-2}{k} \cdot \left(\frac{k}{n-2}\right)^k \cdot \left(\frac{n-2-k}{n-2}\right)^{n-2-k} \\
&\leq \frac{\sqrt{n-2}}{\sqrt{k}\sqrt{n-2-k}}.
\end{aligned}$$

The last inequality follows from Stirling-type upper and lower bounds for factorials, see for example **[Hum40]**. Therefore, we have that Eq. (5.3) is at most

$$m \cdot \min(k-1, n-k-1) + \frac{n^2}{m} \frac{\sqrt{n-2}}{\sqrt{k}\sqrt{n-2-k}}$$

and choosing $m = \Theta\left(\frac{n^{5/4}}{k^{1/4}(n-k)^{1/4}\sqrt{\min(k, n-k)}}\right)$ makes the above quantity $O\left(\frac{n^{7/4}}{k^{1/4}(n-k)^{1/4}}\right)$. Since we could repeat the argument after rotating the plane 180 degrees, the same bound applies to the expected number of k -edges. \square

When $k = \lfloor cn \rfloor$ for some $c \in (0, 1)$, Theorem 5.1.2 shows that the expected number of k -edges of a sample of n points is $O(n^{5/4})$. An artifact of the proof is that when k grows much slower than n , the bound we obtain on the expected number of k -edges is worse. For example, when k is a constant and n grows, the theorem only tells us that the expected number of k -edges is $O(n^{3/2})$. However, we can combine Theorem 5.1.2 with the best known bound for the deterministic k -edge problem to obtain a bound on the expected number of k -edges that gives a uniform bound for all values of k . Indeed if P is any distribution on \mathbb{R}^2 such that the measure of every line is zero, then by Theorem 5.1.2 and the main result of **[Dey97]**, the expected number of k -edges of a sample of n points from P is at most $\min(C_1 n k^{1/3}, \frac{C_2 n^{7/4}}{k^{1/4}(n-2-k)^{1/4}})$ for some constants C_1, C_2 . This quantity is $O(n^{9/7})$ when $k \leq (n-2)/2$.

5.3. On the number of k -edges via the polynomial method

In this section we give another proof of Theorem 5.1.2. The new proof partitions the plane using algebraic curves instead of vertical lines. Given a distribution on \mathbb{R}^2 which has a density, we use the continuous polynomial partitioning theorem of [Gut16b] (Theorem 2.3.4) to obtain an algebraic curve which divides the plane into a number of *cells* of equal probability. The rest of the proof is nearly the same as the proof in Section 5.2.2. The reason this alternative proof is interesting is because it motivates the following open question which, if resolved, may lead to an improvement to the bound in Theorem 5.1.2.

Question 5.3.1. *What is the maximum (finite) number of times that an irreducible non-singular² degree r algebraic curve can intersect the k -edge graph of a set of n points in the plane?*

It is clear that the quantity in Question 5.3.1 is $\Omega(nr)$, and we have some reason to believe that it may be $\Theta(nr)$. The best bound we are able to prove is $O(nr^2)$ (Lemma 5.3.8). This bound is good enough to reprove Theorem 5.1.2 using polynomial partitioning in the case where the distribution has a density (Theorem 5.3.2). Any improvement to our $O(nr^2)$ bound would lead to an improvement in the bound in Theorem 5.3.2: When k is proportional to n , the bound in Theorem 5.3.2 is $O(n^{5/4})$. An $O(nr)$ bound on the quantity in Question 5.3.1 would allow one to improve the bound in Theorem 5.3.2 from $O(n^{5/4})$ to $O(n^{7/6})$ when k is proportional to n .

An $O(nr)$ bound on the quantity in Question 5.3.1 would also have an interesting application to the deterministic k -set problem: It would give another proof of Dey's $O(nk^{1/3})$ bound [Dey97, Dey98] on the maximum number of k -edges of a set of n points in the plane in the case where k is proportional to n . The idea of the proof is as follows. Given any set S of n points in general position in the plane, for some r to be chosen later, use the discrete polynomial partitioning theorem (Theorem 2.3.2) to find a degree $O(r)$ polynomial f such that $\mathbb{R}^2 \setminus Z(f)$ is the union of r^2 pairwise disjoint open sets (called cells) each of which contains at most n/r^2 points of S . The $O(nr)$ bound on the quantity in Question 5.3.1 implies that the number of k -edges formed by two points of S which are both in different cells of the partition is $O(nr)$. Also, the number of k -edges formed by two points which are both in the same cell is at most $r^2 \cdot \binom{n/r^2}{2} = O\left(\frac{n^2}{r^2}\right)$. Since we can assume

²One could consider the same question for possibly singular curves, but, for our purposes, it suffices to consider non-singular curves.

that S is in general position with respect to degree $O(r)$ algebraic curves, we can assume that the number of points contained in $Z(f)$ is $O(r^2)$ and so the number of k -edges formed by two points both of which are in $Z(f)$ is $O(r^4)$. Finally, it is not hard to show that the number of k -edges formed by two points where one point is in $Z(f)$ and the other is not and the interior of the k -edge does not intersect $Z(f)$ is $O(n)$. Indeed, the only way this can happen if one point is in $Z(f)$ and the other is in one of the two cells which have the part of $Z(f)$ that contains the first point on their boundary. So there is a $O(r^2) \cdot 2 \cdot n/r^2$ bound. Moreover, the number of k -edges formed by two points where one point is in $Z(f)$ and the other is not and the interior of the k -edge *does* intersect $Z(f)$ is $O(nr)$, again by the $O(nr)$ bound on Question 5.3.1. Choosing $r = \Theta(n^{1/3})$ shows that the total number of k -edges is $O(n^{4/3})$. See Section 7.4 for a discussion of why we believe the quantity in Question 5.3.1 may be $\Theta(nr)$.

There is one technical issue introduced by our application of the polynomial partitioning theorem to the probabilistic version of the k -edge problem: It can only be applied to distributions which have a density. For this reason, the theorem we prove in this section is slightly weaker than Theorem 5.1.2 because it only applies to distributions which have a density. Although the following theorem is simply a restatement of Theorem 5.1.2 with the added assumption that the distribution has a density, we give a formal statement of the theorem for the sake of readability:

Theorem 5.3.2. *Let P be a probability distribution on \mathbb{R}^2 which has a density. Then for any $1 \leq k \leq n - 1$, $E_P(k, n) = O\left(\frac{n^{7/4}}{k^{1/4}(n-k)^{1/4}}\right)$ (where the constants in big- O are universal).*

Before proving Theorem 5.3.2, in Section 5.3.1 we establish some necessary lemmas concerning algebraic curves.

We make one remark on the requirement that the probability distribution P in Theorem 5.3.2 has a density. As mentioned earlier, in [BS94], Bárány and Steiger construct a probability distribution P with $E_P(\frac{n-2}{2}, n) = \Omega(n \log n)$. The distribution P does not have a density. However, if m_i is any decreasing sequence whose limit is zero, a slight modification to the previously mentioned construction allows Bárány and Steiger to construct a probability distribution P' which has a density and with $E_{P'}(\frac{n-2}{2}, n) = \Omega(m_n n \log n)$. In particular, $E_{P'}(\frac{n-2}{2}, n)$ can still be super-linear even if P' has a density. This shows that the class of distributions which have a density is an important

class of distributions to investigate in the context of the probabilistic k -facet problem. Although the details haven't been worked out as far as we know, it is also probably not hard to construct a distribution P which has a density and with $E_P(k, n) = \Omega(m_n n e^{\Omega(\sqrt{\log n})})$

5.3.1. Counting intersection points of an algebraic curve and the k -edge graph. Any polynomial $f \in \mathbb{R}[x_1, x_2]$ defines an algebraic curve $Z(f) := \{x \in \mathbb{R}^2 : f(x) = 0\}$.

Our use of the polynomial partitioning technique requires us to bound the number of times the k -edge graph G_k of a set of n points can intersect a degree r algebraic curve $Z(f)$, i.e., we must give some answer to Question 5.3.1. Note that the number of points of intersection of G_k and $Z(f)$ could be infinite if $Z(f)$ contained one of the lines spanned by a k -edge in G_k . However, for our purposes, it suffices to bound the number of intersection points in the case when it is finite.

In order to establish our $O(nr^2)$ bound on the quantity in Question 5.3.1, we first show how to partition an irreducible algebraic curve into the union of $O(r^2)$ convex and concave pieces (Proposition 5.3.6). Combining the convex/concave chains decomposition of the k -edge graph G_k with the partition of a degree r algebraic curve into $O(r^2)$ convex and concave pieces allows us to show that a degree r algebraic curve intersects the k -edge graph of a set of n points at most $O(nr^2)$ times assuming the number of intersections is finite (Lemma 5.3.8).

First, we show how to partition an irreducible curve $Z(f)$ into the union of a finite number of points and a finite number of convex/concave x -monotone connected curves.

Definition 5.3.3. *A connected curve $C \subset \mathbb{R}^2$ is x -monotone if every vertical line intersects it in at most one point.*

Definition 5.3.4. *An x -monotone curve C is convex (respectively, concave) if for every three points $(x_1, y_1), (x_2, y_2), (x_3, y_3) \in C$ with $x_1 < x_2 < x_3$, the point (x_2, y_2) is below (respectively, above) or on the line joining (x_1, y_1) and (x_3, y_3) .*

In order to break $Z(f)$ into convex/concave pieces, we need to use the inflection points of $Z(f)$.

Definition 5.3.5 ([Kir92]). *A non-singular point (a, b) of an algebraic curve $Z(f)$ is an inflection point if the Hessian curve $H_f(x, y) := f_y^2 f_{xx} - 2f_x f_y f_{xy} + f_x^2 f_{yy}$ is equal to zero at (a, b) . (The notation f_x denotes the partial derivative with respect to x .)*

Proposition 5.3.6. *A non-singular irreducible curve $Z(f) \subset \mathbb{R}^2$ of degree r that is not a vertical line can be partitioned into the union of at most $4r^2$ points and at most $6r^2$ x -monotone curves where each x -monotone curve is either convex or concave.*

PROOF. If $Z(f)$ is a non-vertical line, the conclusion is clearly true. So assume that $Z(f)$ is not a line. Let $F = Z(f) \cap Z(f_y)$. We know that f depends on y and not just x because otherwise $Z(f)$ would be a vertical line. This means that f_y is not identically zero. Now because f is irreducible and the degree of f_y is less than the degree of f , the polynomials f and f_y cannot have a common factor. Therefore, by Bézout's theorem, $|F| \leq r(r-1)$.

Let I be the set of inflection points of $Z(f)$. An irreducible curve of degree $r \geq 2$ has at most $3r(r-2)$ inflection points [Kir92, Proposition 3.33] so $|I| \leq 3r(r-2)$. Let \mathcal{C} be the set of connected components of $Z(f) \setminus (I \cup F)$. Because of the removal of the points in F , every curve in \mathcal{C} is x -monotone. To show this, let $C \in \mathcal{C}$. Assume there exists two points $a, b \in C$ that have the same x -coordinate. Since C is not a vertical line, there must be a point $x \in C$ that is between a, b and such that x is not on the line through a, b . Therefore, between a and b , the curve must travel in the positive x direction and then in the negative x -direction, meaning there exists a point on the curve between a and b where $f_y = 0$, a contradiction.

Now we show that, because of the removal of the inflection points I , every curve in \mathcal{C} is either a convex x -monotone curve or a concave x -monotone curve. Each curve in \mathcal{C} is the graph of a function defined in an interval. We claim that, for each curve in \mathcal{C} , the second derivative of the associated function exists everywhere and is never zero. Let $C \in \mathcal{C}$. Since C does not contain a point where $f_y = 0$, using the implicit function theorem, for each $(u, v) \in C$, there exists a smooth function $\phi : (u - \epsilon, u + \epsilon) \rightarrow \mathbb{R}$ that gives a local parameterization of the curve near (u, v) [Rut00, Theorem 4.22]. Now a simple calculation shows that if $\phi''(x)$ is equal to 0 at x , then the Hessian curve $f_y^2 f_{xx} - 2f_x f_y f_{xy} + f_x^2 f_{yy}$ equals zero at $(x, \phi(x))$. Indeed, we have parameterized $Z(f)$ near a given point in C by $x \mapsto (x, \phi(x))$. Differentiating $f(x, \phi(x)) = 0$ gives $f_x + \phi'(x)f_y = 0$ and so $\phi'(x) = -\frac{f_x}{f_y}$. Differentiating again gives $\phi''(x)f_y + (1, \phi'(x)) \begin{pmatrix} f_{xx} & f_{xy} \\ f_{xy} & f_{yy} \end{pmatrix} \begin{pmatrix} 1 \\ \phi'(x) \end{pmatrix} = 0$ and now rearranging and using the fact that $\phi'(x) = -\frac{f_x}{f_y}$ and the fact that $f_y \neq 0$ shows that if $\phi''(x) = 0$ then the Hessian curve is zero. The inflection points of $Z(f)$ are precisely the points where the

Hessian curve is zero. Therefore, since all inflection points were removed, the second derivative of the function whose graph is C is never zero. This means that the function is either strictly convex or strictly concave, and so C is either a convex or concave x -monotone curve.

Now we determine how many distinct curves \mathcal{C} can contain. The number of connected components of $Z(f)$ is at most $2r^2$ by either [Mil64, Theorem 2] or [She20, Theorem 2.7]. We removed at most $r(r-1) + 3r(r-2)$ points from $Z(f)$. Because $Z(f)$ is non-singular, it has no points of self-intersection. Therefore, each point which is removed increases the number of connected components of $Z(f) \setminus (I \cup F)$ by at most 1. Therefore, the number of connected components of $Z(f) \setminus (I \cup F)$ is at most $2r^2 + r(r-1) + 3r(r-2) \leq 6r^2$. \square

The decomposition into convex/concave pieces is useful because of the following fact:

Lemma 5.3.7. *Let C be an x -monotone convex curve and D an x -monotone concave curve. If the number of points of intersection of C and D is finite, then it is at most 2.*

PROOF. Assume that C and D intersect in three points $(x_1, y_1), (x_2, y_2), (x_3, y_3)$. Observe that the three points $(x_1, y_1), (x_2, y_2), (x_3, y_3)$ must be contained in a line ℓ and we may assume that $x_1 < x_2 < x_3$. We claim that C and D must both contain the line segment connecting the three points. Indeed, assume that C does not contain the line segment connecting (x_1, y_1) and (x_2, y_2) . Then there must be a point $(x_0, y_0) \in C$ with $x_1 < x_0 < x_2$ and (x_0, y_0) strictly below the line connecting (x_1, y_1) and (x_2, y_2) . But then the point (x_2, y_2) is above the line connecting (x_0, y_0) and (x_3, y_3) , a contradiction to convexity of C . The argument for the other cases is similar. \square

Now we can establish the bound on the number of intersection points between $Z(f)$ and G_k .

Lemma 5.3.8. *Let $S \subset \mathbb{R}^2$ be a set of points in general position, $G_k = (S, E_k)$ the k -edge graph of S , and $f \in \mathbb{R}[x_1, x_2]$ a degree r polynomial such that all irreducible components of $Z(f)$ are non-singular and $S \cap Z(f) = \emptyset$. If the number of points of intersection of $Z(f)$ and G_k is finite, then it is at most $13nr^2$.*

PROOF. First assume that $Z(f)$ is irreducible. If $Z(f)$ is a line, then it follows from Lemma 5.2.2 that the number of intersection points is at most $\max(k-1, n-k+1) \leq 13n$ and we are done. So assume that $Z(f)$ is not a line. First, we need to decompose the curve $Z(f)$ into the union of

convex and concave pieces. By Proposition 5.3.6, $Z(f)$ can be partitioned into the union of $6r^2$ convex/concave x -monotone curves and at most $4r^2$ points. Let A be the set of convex x -monotone curves, B the set of concave x -monotone curves, and N the set of points in the partition.

By Lemma 5.2.2, G_k can be decomposed into the union of $k - 1$ convex chains C_1, \dots, C_{k-1} , or $n - k + 1$ concave chains D_1, \dots, D_{n-k+1} .

Recall that we are assuming that no line spanned by two points from S is vertical. Therefore, the convex chains C_1, \dots, C_{k-1} , and the concave chains D_1, \dots, D_{n-k+1} never contain two points on a vertical line. Furthermore, we claim that any convex or concave chain C_i or D_j intersects $Z(f)$ in only finitely many points. Indeed, if one of these chains intersected $Z(f)$ in infinitely many points, one of the line segments in the chain would have to intersect $Z(f)$ in infinitely many points. Recall the fact that if a degree r algebraic curve intersects a line in more than r points, then the curve must contain the line. Since we are assuming that $Z(f)$ does not contain any of the points in S , this is not possible. Thus, we can apply Lemma 5.3.7 and the concave chain decomposition of G_k to show that the number of k -edges in G_k that intersect $Z(f)$ at a point contained in one of the convex x -monotone curves in A is at most $2(n - k + 1)6r^2$. Similarly, the number of k -edges in G_k that intersect $Z(f)$ at a point contained in one of the concave x -monotone curves in B is at most $2(k - 1)6r^2$. Additionally, there are $4r^2$ points in the set $N \subset Z(f)$ which are not contained in any of the convex/concave x -monotone curves. Therefore, the total number of intersections is at most $2(n - k + 1)6r^2 + 2(k - 1)6r^2 + 4r^2 \leq 13nr^2$.

If $Z(f)$ is not irreducible, then say $Z(f)$ is the union of m irreducible components f_1, f_2, \dots, f_m of degrees r_1, r_2, \dots, r_m . By the above, the number of intersection points of $Z(f_i)$ and G_k is at most $13nr_i^2$. So the total number of intersection points of $Z(f)$ and G_k is at most $\sum_{i=1}^m 13nr_i^2 \leq 13nr^2$. \square

5.3.2. Proof of Theorem 5.3.2.

PROOF. Let W be the density of P . For any fixed n , we use Theorem 2.3.4 applied to W to find a degree $r := r(n)$ (to be chosen later) polynomial f which divides $\mathbb{R}^2 \setminus Z(f)$ into a family \mathcal{O} of $\Theta(r^2)$ pairwise disjoint open sets such that for all $O \in \mathcal{O}$, the integrals $\int_O W$ are within a factor of 2 of each other. Furthermore, all irreducible components of $Z(f)$ are non-singular.

Let $X = \{X_1, \dots, X_n\}$ be a sample of n points from P . Observe that the probability that two points in X span a vertical line is 0 so we do not need to consider this case. Also, since P has a density, the measure of every line is 0. This means that X is in general position with probability 1 so we may assume this as well. Therefore, we can analyze the k -edges of X using the k -edge graph G_k of X . Also, since the Lebesgue measure of $Z(f)$ is zero, $X \cap Z(f) = \emptyset$ with probability 1 so we assume this as well.

We compute the expected number of k -edges in G_k by considering two different types of k -edges separately. First we bound the expected number of k -edges formed by two points in different cells of the partition. Then we bound the expected number of k -edges formed by two points in the same cell of the partition. That is, the expected number of k -edges in G_k is equal to

$$(5.6) \quad \begin{aligned} & \mathbb{E}(\text{number of } k\text{-edges in } G_k \text{ formed by two points in different cells}) \\ & + \mathbb{E}(\text{number of } k\text{-edges in } G_k \text{ formed by two points in the same cell}). \end{aligned}$$

If $\text{conv}(X_1, X_2)$ is a k -edge formed by two points X_1, X_2 in different cells, then $\text{conv}(X_1, X_2)$ intersects $Z(f)$. So to bound the expected number of k -edges in G_k formed by two points in different cells, it suffices to bound the expected number of k -edges that intersect $Z(f)$. We claim that the number of points of intersection between G_k and $Z(f)$ is finite. This is true because otherwise some k -edge would have to intersect $Z(f)$ infinitely many times. If a degree r algebraic curve intersects a line more than r times it must contain that line. If this were true, then $Z(f)$ would have to contain the two points of X forming the line, but we are assuming that $X \cap Z(f) = \emptyset$. Therefore, the number of points of intersection between G_k and $Z(f)$ is finite and so we can apply Lemma 5.3.8 to show that the first term in Eq. (5.6) is at most $13nr^2$.

Now we bound the second term. Recall that $Z(f)$ divides $\mathbb{R}^2 \setminus Z(f)$ into a family \mathcal{O} of $\Theta(r^2)$ pairwise disjoint open sets (called cells) such that for all cells $O \in \mathcal{O}$, the integrals $\int_O W$ are within a factor of 2 of each other. Therefore, for any fixed $i \neq j$, the probability that X_i and X_j are in the same cell is at most C/r^2 for a universal constant $C > 0$.

Now, the second term in Eq. (5.6) can be bounded using nearly the same argument which we used to bound the second term in Eq. (5.3) in the proof of Theorem 5.1.2. This argument is given in Eqs. (5.4) and (5.5) in the proof of Theorem 5.1.2. The only change is that for any

$X_i, X_j \in X$, the probability $\mathbb{P}(X_1, X_2 \text{ are in the same cell})$ is now upper bounded by C/r^2 instead of $1/m$. Therefore, we have that Eq. (5.6) is at most

$$13nr^2 + C \frac{n^2}{r^2} \frac{\sqrt{n-2}}{\sqrt{k}\sqrt{n-2-k}}$$

and choosing $r^2 = \Theta\left(\frac{n^{3/4}}{k^{1/4}(n-k)^{1/4}}\right)$ gives $O\left(\frac{n^{7/4}}{k^{1/4}(n-2-k)^{1/4}}\right)$. Since we could repeat the argument after rotating the plane 180 degrees, the same bound applies to the expected number of k -edges. \square

Translations of a fixed convex body in the plane

We study a variation on the k -set problem for the set system whose set of ranges consists of all translations of some strictly convex body in the plane. The motivation is to show that the technique by Bárány and Steiger is tight for a natural family of set systems. For any such set system, we determine bounds for the expected number of k -sets which are tight up to logarithmic factors.

In particular, we show that the two-step argument in the proof of Theorem 5.1.1 (Eq. (5.1) and the general upper bound of the integrand in Eq. (5.2)) is not as loose as it seems, if one applies it to the k -set problem on a set system other than half-planes (generalized in a natural way). More precisely, let $C \subseteq \mathbb{R}^2$ be the interior of a fixed convex body. We consider the set system of translations of C and study the expected number of k -sets and k -edges. In addition, we show some deterministic bounds to put our probabilistic bounds in context.

For the case where C is strictly convex, we show:

- A relation between k -sets and k -edges that allows one to derive upper bounds on the number of k -sets from upper bounds on the number of k -edges (Lemma 6.2.5).
- For certain natural distributions, the expected number of k -sets and k -edges for a random set of n points and some k proportional to n is $\Theta^*(n^{3/2})$ (where $*$ means that polylogarithmic factors are ignored) (Theorem 6.2.9 and Theorem 6.5.3). The upper bound uses the Bárány and Steiger technique, while the lower bound uses the uniform convergence theorem of Vapnik and Chervonenkis [VC71].
- The growth function is $O(n^2)$ (Proposition 6.4.1).

For the case where C has C^2 boundary, we show that the maximum number of k -sets of n points with k proportional to n is $\Omega(n^2)$ (Theorem 6.3.2).

Some of the assumptions above are chosen for readability, the actual theorems have weaker assumptions in some cases.

6.1. Introduction

In this chapter we study a natural variation of the k -set problem for translations of a fixed convex set on the plane, namely, the number of ways in which one can enclose k points out of a given finite set of points by a translation of a convex set so that its boundary strictly separates them from the rest. We will show nearly matching upper and lower bounds on the expected number of ways.

For our lower bound, one of our tools will be the uniform convergence theorem of Vapnik and Chervonenkis [VC71]. This introduces a minor technical complication: their theorem is about abstract set systems without regard to whether sets have a boundary, while the standard k -set problem for lines on the plane ask for strict separation by a line and therefore the natural choice for our generalization is to ask for strict separation by a curve.

Similarly, the other side of our argument, our upper bound, is a variation on the two-step argument in the proof of Theorem 5.1.1 (Eq. (5.1) and the general upper bound of the integrand in Eq. (5.2)), which uses k -edges and therefore also uses the boundary curve in a fundamental way.

A *convex body* is a compact convex set with non-empty interior. A set C is *strictly convex* if for all $x, y \in C$ with $x \neq y$ and for all $\lambda \in (0, 1)$ we have $\lambda x + (1 - \lambda)y \in \text{int } C$. A *set system* is a pair (X, \mathcal{R}) , where X is a set and \mathcal{R} is a family of subsets of X . The elements of \mathcal{R} are called *ranges*. For a set $C \subseteq \mathbb{R}^2$, let (\mathbb{R}^2, T_C) be the set system of translations of C (that is, T_C is the family of translation of C). We are interested in translations of convex sets and it will be notationally convenient to set C to be the *interior* of a fixed convex body. So, for this section, C will be restricted (at least) to be the interior of a convex body. In this case, when we say that a point lies on the boundary of a range, the point does not lie in the range.

Definition 6.1.1. *For a finite subset $S \subseteq \mathbb{R}^2$, a T_C - k -set of S is a subset $T \subseteq S$ of k points such that for some $Q \in T_C$, $S \cap \text{bd } Q = \emptyset$ and $T = S \cap Q$.*

6.2. Upper bound for T_C - k -sets, probabilistic, k proportional to n

This section establishes our upper bound on the expected number of T_C - k -sets of a set of n iid points when C is the interior a strictly convex body and k is proportional to n . So for Section 6.2, let $C \subseteq \mathbb{R}^2$ be the interior of a strictly convex body.

Definition 6.2.1. *A set of points in \mathbb{R}^2 is in general position relative to C if no three points lie on the boundary of some range in T_C (i.e., some translation of C).*

Lemma 6.2.2. *Let (p, q) be a pair of distinct points in \mathbb{R}^2 . Then there are at most two ranges $x + C$ satisfying $p, q \in \text{bd}(x + C)$.*

PROOF. Up to a rotation we can assume that $r := q - p$ is vertical. Suppose for a contradiction that there are three ranges $x + C$ satisfying $p, q \in \text{bd}(x + C)$. This implies there are three points p_1, p_2, p_3 such that $p_1, p_2, p_3, p_1 + r, p_2 + r, p_3 + r \in \text{bd} C$. Let $f(x_1)$ denote the length of segment “ C intersected with the vertical line at $x_1 \in \mathbb{R}$.” Function f is positive in a non-empty interval (a, b) , is strictly concave in $[a, b]$ and takes value $\|r\|$ at three points in $[a, b]$. But there is no such function so this is a contradiction. \square

From the lemma we conclude:

Corollary 6.2.3. *Let $V \subseteq (\mathbb{R}^2)^2$ be the set of pairs of distinct points that can appear on the boundary of some range. Then there exists a continuous onto function $C : V \rightarrow T_C$.*

PROOF. Every pair in V can appear on the boundary of one or two ranges. When a pair appears on the boundary of exactly one range, map both orderings of the pair to that range. When a pair (p, q) appears on the boundary of two ranges, let C_1 and C_2 be the two translations of C containing p and q on their boundaries. Let $C(p, q)$ be the unique translation that solves $\max_{i \in \{1, 2\}} \text{area}(C_i \cap \text{aff}(p, q)^+)$ (where $\text{aff}(p, q)^+ = \{r : -(q-p)_x(r-p)_y + (q-p)_y(r-p)_x > 0\}$). \square

From now on we let $C(\cdot, \cdot)$ denote the function given by Corollary 6.2.3 (with a slight abuse of notation).

Definition 6.2.4. For a set of points S in general position relative to C , let a (oriented) T_C - k -edge be an ordered pair of points $(p, q) \in V$ with $p, q \in S$ ($p \neq q$) such that $C(p, q)$ contains k points of S .

We now show a bound relating T_C - k -sets and T_C - k -edges, which follows from a variation of known continuous deformation arguments [HP11, Lemma 5.15], [Mat02, Chapter 11]. For a finite set of points S in general position relative to C , let $e_k(S)$ be the number of T_C - k -edges of S and let $a_k(S)$ be the number of T_C - k -sets of S .

Lemma 6.2.5. For a finite set of points S in general position relative to C and $k \geq 2$ we have $a_k(S) \leq 4(e_{k-2}(S) + e_{k-1}(S) + e_k(S))$.

PROOF. To prove the claim we will construct an injective function f from T_C - k -sets of S to a labelled extension of the set of T_C - k -edges. The function is defined as follows: Let $Q \subseteq S$ be a T_C - k -set induced by some range C_0 . Translate C_0 in the x direction until some point $p \in S$ lies on its boundary to obtain range C' , then translate C' while keeping p on its boundary (letting p slide along the boundary) until another point $q \in S$ lies on the boundary to obtain a range C'' (there may be more than one choice here, pick arbitrarily). We have that $\text{cl } C''$ contains Q and between zero and two other points from S . From the general position assumption, $\text{bd } C''$ contains exactly two points from S . Swap points p, q if needed so that $C'' = C(p, q)$. Pick labels $l_p, l_q \in \{IN, OUT\}$ for p and q according to whether they are in Q . This completes the definition of an f from T_C - k -sets of S to $S^2 \times \{IN, OUT\}^2$, namely $f(Q) = (p, q, l_p, l_q)$.

We now show that f is injective. Let Q, Q' be two T_C - k -sets of S induced by ranges C_0, C'_0 , respectively, and so that $f(Q) = f(Q') = (p, q, l_p, l_q)$. By definition of f we have that Q is equal to $S \cap C(p, q)$ with p and q added according to the labels. But then by definition of f we have that Q' is also equal to that set and therefore equal to Q . This establishes that f is injective.

To conclude, the image of f contains only pairs (p, q) that are T_C - r -edges of S for $r \in \{k - 2, k - 1, k\}$. The claim follows. \square

Assumption 6.2.6. Given C , probability distribution P on \mathbb{R}^2 is such that $P(\text{bd}(x + C)) = 0$ for all $x \in \mathbb{R}^2$.

(In particular the assumption on P holds if P has a density.)

Proposition 6.2.7. *Let P be a Borel probability distribution satisfying Assumption 6.2.6. Let Y, Z be a pair of iid points, each according to P . Then $Y \neq Z$ a.s.*

PROOF. Fix a point b on the boundary of C (so that the origin is on the boundary of $-b + C$). Note that $\mathbb{P}(Y = Z) = \mathbb{E}(\mathbb{P}(Y = Z \mid Z)) \leq \mathbb{E}(\mathbb{P}(Y \in \text{bd}(Z - b + C) \mid Z)) = 0$. \square

Proposition 6.2.8. *Let P be a Borel probability distribution satisfying Assumption 6.2.6. Let X be a random set of n iid points, each according to P . Then X is in general position relative to C a.s.*

PROOF. It is enough to prove the claim for $n = 3$. Let Y, Z, W be three iid random points according to P . By Proposition 6.2.7, $Y \neq Z$ a.s. Then

$$\begin{aligned} & \mathbb{P}((\exists a)Y, Z, W \in \text{bd}(a + C)) \\ &= \mathbb{P}(Y \neq Z, (\exists a)Y, Z, W \in \text{bd}(a + C)) \\ &= \mathbb{E}\left(\mathbb{P}((\exists a)Y, Z, W \in \text{bd}(a + C) \mid Y, Z) \mid Y \neq Z\right) \\ &= \mathbb{E}\left(\mathbb{P}(W \in \text{bd}C(Y, Z) \text{ or } W \in \text{bd}C(Z, Y) \mid Y, Z) \mid Y \neq Z\right) \\ &= 0. \end{aligned}$$

\square

Theorem 6.2.9 (T_C - k -set/edge upper bound, probabilistic, k proportional to n). *Let $c \in (0, 1)$. Let P be a Borel probability distribution satisfying Assumption 6.2.6. Let X be a random set of n iid points, each according to P . Let A_n (resp. E_n) be the number of T_C - k -sets (resp. T_C - k -edges) of X for $k = \lfloor cn \rfloor$. Then*

$$\mathbb{E}(E_n) \leq O(n^{3/2})$$

and

$$\mathbb{E}(A_n) \leq O(n^{3/2})$$

(where the constants in big- O depend only on c).

PROOF. Let $C(p, q)$ and V be as in Corollary 6.2.3. From Proposition 6.2.8, X is in general position relative to C a.s.

Let $X = \{X_1, \dots, X_n\}$. Let $T = P(C(X_1, X_2))$ with the additional convention that $C(p, q) = \emptyset$ if $(p, q) \notin V$. In this way T is defined whenever $X_1 \neq X_2$, that is, a.s. by Proposition 6.2.7. Let $G(t) = \mathbb{P}(T \leq t)$ for $t \in \mathbb{R}$. Using a variation of Eq. (5.1) and the argument in the proof of Theorem 5.1.2 we get:

$$\begin{aligned}
& \mathbb{P}((X_1, X_2) \text{ is a } T_C\text{-}k\text{-edge of } X) \\
&= \mathbb{P}(|C(X_1, X_2) \cap (X \setminus \{X_1, X_2\})| = k) \\
&= \mathbb{E}\left(\mathbb{P}(|C(X_1, X_2) \cap (X \setminus \{X_1, X_2\})| = k \mid X_1, X_2)\right) \\
&= \binom{n-2}{k} \mathbb{E}(T^k(1-T)^{n-2-k}) \\
&= \binom{n-2}{k} \int_0^1 t^k(1-t)^{n-2-k} dG(t)
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}(E_n) &= n(n-1) \binom{n-2}{k} \int_0^1 t^k(1-t)^{n-2-k} dG(t) \\
&\leq n^2 \binom{n-2}{k} \left(\frac{k}{n-2}\right)^k \left(\frac{n-2-k}{n-2}\right)^{n-2-k} \\
&\leq n^2 \frac{\sqrt{n-2}}{\sqrt{k}\sqrt{n-2-k}} \\
&\leq O(n^{3/2}).
\end{aligned}$$

This proves the first inequality. From this, the second inequality is immediate using Lemma 6.2.5. □

6.3. Lower bound for T_C - k -sets, deterministic, k proportional to n

In this section we show a lower bound on the maximum number of T_C - k -sets of a broad family of set systems of the form (\mathbb{R}^2, T_C) for k proportional to n . We illustrate the main idea of the argument in Proposition 6.3.1 for the case where C is a unit square. In Theorem 6.3.2, we then use

the argument for the case where C is the interior of a convex body with C^2 boundary (actually, slightly more general than that).

While the sets of points in the following results may not be in general position, this is not a true weakness of the results. The reason is that, like in the case of the standard k -set problem for lines, the number of T_C - k -sets of a given set of points cannot decrease by applying any sufficiently small perturbation to the points, because any range inducing a T_C - k -set must by definition contain no point on its boundary. Thus, the maximum number of T_C - k -sets among set in general position is no smaller than the number of T_C - k -sets among unrestricted sets of points.

Proposition 6.3.1 (idea of lower bound for maximum, the cross). *Let C be the open unit square. Let $0 < c < c' < 1$. Then for $cn \leq k \leq c'n$ we have $\max_{|S|=n} a_k(S) = \Omega(n^2)$ (where the constants in Ω depend only on c and c').*

PROOF. We show the case where n is a multiple of 4 and $k = n/2$, the rest is similar. Consider a set of points equally spaced on the x and y axes forming a cross. Say, for $\lambda = 4/n$, let $S = \lambda(\mathbb{Z}^2 \cap (x\text{-axis} \cup y\text{-axis}) \cap [-n/4, n/4]^2 \setminus (0, 0))$. Then $|S| = n$ and the claim follows. \square

For the next result we assume that the boundary of C is well approximated by its unique tangent line at certain points (locally of class C^2). See [Gru07, Section 5.1, subsection ‘‘Second-Order Differentiability’’] for basic facts about differentiability of the boundary of a convex body.

Theorem 6.3.2 (lower bound for maximum, the cross). *Assume that $C \subseteq \mathbb{R}^2$ is the interior of a convex body such that there exist linearly independent unit vectors $u, v \in \mathbb{R}^2$ and points $a, b, c, d \in \text{bd } C$ such that $\text{bd } C$ is C^2 in a neighborhood of a, b, c, d with outer normals $u, v, -u, -v$, resp. Let $0 < c < c' < 1$. Then for $cn \leq k \leq c'n$ we have $\max_{|S|=n} a_k(S) = \Omega(n^2)$ (where the constants in Ω depend only on c, c' and C).*

PROOF. We show the case where n is a multiple of 8 and $k = n/2$, the rest follows easily. Up to an invertible linear transformation, we can assume $u = e_1$ and $v = e_2$, without loss of generality. Let $U = \{e_1, e_2, -e_1, -e_2\}$ and for $p \in U$ let $v(p)$ be a locally C^2 point on the boundary of C and having outer normal p .

For $p \in U$ and some $t > 0$, consider the segment of length $2t$ perpendicular to the boundary of C at $v(p)$ and centered at $v(p)$, namely $s(p) := \text{conv}\{v(p) - tp, v(p) + tp\}$. Finally, consider a

(one-dimensional) grid $g(p)$ of $n/4$ equally spaced points on segment $s(p)$. Let our set of points be $S = \cup_{p \in U} g(p)$. By construction $|S \cap C| = n/2$. Let $\epsilon := 2t/(n/4 - 1)$ be the gap between consecutive points in each segment.

We will now show that we can choose $t > 0$ small enough so that there are $\Omega(n^2)$ small translation of C that induce different subsets of S , each containing $n/2$ points. The idea of the argument is to translate C independently in the vertical and horizontal direction, to pick $\Omega(n)$ different subsets of $n/4$ points among the pair of vertical segments and similarly for the horizontal segments. The translations, notated $p + C$ and parameterized by p , form the following grid around the origin: $G := \{(k\epsilon, l\epsilon) \in \mathbb{R}^2 : k, l \in \mathbb{Z} \cap [-n/8, n/8]\}$. By Taylor's theorem and compactness there exist constants $\alpha > 0, t_M > 0$ (that depend only on C) such that the boundary of C has a C^2 parametrization $y = \phi(x)$ in a neighborhood of $(x_0, y_0) := v(-e_2)$ such that $|\phi(x) - y_0| \leq \alpha(x - x_0)^2$ for $x \in [x_0 - t_M, x_0 + t_M]$ (and similarly for $v(e_1), v(-e_1), v(e_2)$). In particular, $|\phi(x) - y_0| \leq \alpha t_M^2$. We choose $t > 0$ small enough so that C contains the same subset of $g(-e_2)$ when translated distance less than or equal to t in the horizontal direction. Note that this also ensures that the boundaries of those translations contain no point of $g(-e_2)$. The nearest point from $g(-e_2)$ to the line $y = y_0$ is at distance $\epsilon/2 = \frac{t}{n/4-1} > 4t/n$ so it is enough to have $\alpha t^2 \leq 4t/n$, that is, we set $t = \min\{t_M, \frac{4}{\alpha n}\}$.

With these choices, every $p \in G$ induces a different T_C - k -set of S with $k = n/2$ and therefore $a_{n/2}(S) \geq |G| \geq n^2/16$. □

To understand the scope of Theorem 6.3.2, note that the condition on C is satisfied when C is the interior of a convex body with C^2 boundary. Also, no triangle satisfies the assumptions of Theorem 6.3.2.

6.4. Bounds on the growth function

To put our results on the number of T_C - k -sets and T_C - k -edges in context, we state some basic bounds on the growth function of set system (\mathbb{R}^2, T_C) , namely the maximum number of subsets of a set of n point in \mathbb{R}^2 induced by translations of C . For simplicity some of our bounds have extra assumptions on C that may not be necessary.

The *growth function* [VC71] of set system (\mathbb{R}^2, T_C) is given by

$$n \mapsto \max_{S \subseteq \mathbb{R}^2, |S|=n} |\{S \cap (x + C) : x \in \mathbb{R}^2\}|.$$

Proposition 6.4.1. *Let $C \subseteq \mathbb{R}^2$ be the interior of a strictly convex body. The growth function of (\mathbb{R}^2, T_C) is at most $n^2 - n + 2$.*

PROOF. For the proof we use the notions of a dual set system and dual growth function.

The first step is to notice that (\mathbb{R}^2, T_C) corresponds to the dual set system of (\mathbb{R}^2, T_{-C}) : a range $x - C \in T_{-C}$ corresponds to point $x \in \mathbb{R}^2$ and a point $y \in \mathbb{R}^2$ corresponds to range $y + C$, together with the equivalence “ $y \in x - C$ is equivalent to $x \in y + C$ ”. In this way, the growth function of (\mathbb{R}^2, T_C) is the dual growth function of (\mathbb{R}^2, T_{-C}) .

The second step is to bound the dual growth function of (\mathbb{R}^2, T_{-C}) . Its value at n is bounded by the number of connected components of the complement of n translations of $\text{bd}(-C)$, or, equivalently, $\text{bd}(C)$. Adding n translations of C one by one, this number of connected components is 2 for $n = 1$ and, using the fact that two translations of $\text{bd}(C)$ intersect in at most two points, it increases by at most $2(k-1)$ when the k th translation is added. Therefore the number of connected components is at most $n^2 - n + 2$. \square

For clarity we state the following summarizing result:

Theorem 6.4.2. *Let $C \subseteq \mathbb{R}^2$ be the interior of a strictly convex body with C^2 boundary. The growth function of (\mathbb{R}^2, T_C) is $\Theta(n^2)$ (where the constants in Θ depend only on C).*

PROOF. Immediate from Proposition 6.4.1 and Theorem 6.3.2. \square

In order to prove our results in Section 6.5 in more generality, we state here a weaker bound on the growth function with weaker assumptions on C .

Theorem 6.4.3. *Let $C \subseteq \mathbb{R}^2$ be the interior of a convex body. The VC-dimension of (\mathbb{R}^2, T_C) is at most 3. The growth function of (\mathbb{R}^2, T_C) is at most $\binom{n}{0} + \binom{n}{1} + \binom{n}{2} + \binom{n}{3} \leq (en/3)^3$.*

PROOF. A special case of a result in [NT10] establishes that the VC-dimension of (\mathbb{R}^2, T_C) is at most 3 when $C \subseteq \mathbb{R}^2$ is a convex body. The bound extends to our case (interior of a convex

body) using the observation that if translations of the interior a convex body C shatter a given finite set of points then translations of a scaled down C also shatter the same set. The rest follows from the Sauer-Shelah lemma. \square

The VC-dimension bound is tight: C equal to any fixed triangle is a tight example.

6.5. Lower bound for T_C - k -sets, probabilistic, some k proportional to n

In this section we show, for some k proportional to n , an $\Omega^*(n^{3/2})$ lower bound for the expected number of T_C - k -sets for a random sample of n points from the uniform distribution in a set $A \subseteq \mathbb{R}^2$ sufficiently large to contain translations of C . The restriction to a subset A is necessary as there is no uniform distribution in \mathbb{R}^2 . Our argument uses crucially the fact that translations of C that are contained in A have the same probability under the uniform distribution in A . The minor technical complications introduced by the fact that A is bounded could be avoided by considering a similar set system of translations of a disk (say) on the surface of the two-dimensional sphere (or translations of a shape on the flat torus) with the uniform distribution (a version not studied in this paper).

The idea of the proof is the following: First show that for a random sample X of n points in A , with high probability the number of induced subsets by translations of C contained in A is $\Omega^*(n^2)$ (Lemma 6.5.2). Then, by VC's uniform convergence theorem, with high probability each translation of C contained in A contains $cn + O^*(\sqrt{n})$ points from X for some c . Therefore, by the pigeonhole principle there are $\Omega^*(n^{3/2})$ induced subsets of X that contain exactly the same number of points, that is, X has $\Omega^*(n^{3/2})$ T_C - k -sets for some k .

We start by showing that if two translations of C are close then they have a large intersection.

Lemma 6.5.1. *Let $C \subseteq \mathbb{R}^2$ be the interior of a convex body that contains a unit ball. If $\|x\| \leq 1$, then*

$$\text{area}(C \cap (x + C)) \leq (1 - \|x\|/2) \text{area}(C).$$

PROOF. Consider the function $f(x) = \text{area}(C \cap (x + C))$. It is logconcave (by the Prékopa-Leindler inequality and the fact that $f(x) = \mathbf{1}_C(x) * \mathbf{1}_{-C}(x)$). Also, $f(0) = \text{area}(C) \geq \pi$. In other words, $\log f(x)$ is concave and, while it is not differentiable at $x = 0$, we can use directional

derivatives and tangent rays at $x = 0$ to upper bound it by a function of the form $x \mapsto \log f(0) + c\|x\|$, where $c < 0$ is an upper bound on the one-sided directional derivative.

We calculate a suitable c now. The one-sided directional derivative at 0 along unit vector $v \in \mathbb{R}^2$ is

$$Df(0)(v) = -2 \text{length}(\text{projection of } C \text{ onto line perpendicular to } v) \leq -4$$

(from the analysis of the movement of chords of C parallel to v : as a chord moves by distance Δt , it contributes $2\Delta t$ and the family of chords is parameterized by values in projection of C onto line perpendicular to v). Thus, $\log f(0) = \log(\text{area } C)$ and $D(\log f)(0)(v) = Df(0)(v)/f(0) \leq -4/\pi \leq -1$ (i.e. we can take $c = -1$) and these estimates with concavity of $\log f(x)$ give $\log f(x) \leq \log(\text{area}(C)) - \|x\|$. That is, $f(x) \leq \text{area}(C)e^{-\|x\|}$. We use the inequality $e^{-t} \leq 1 - t(1 - 1/e)$ for $t \in [0, 1]$ to conclude that if $\|x\| \leq 1$, then $f(x) \leq \text{area}(C)(1 - \|x\|(1 - 1/e))$. The claim follows. \square

We show now that the number of ranges induced by translations of C on certain random sets of n points is $\Omega^*(n^2)$. Because this is meant to be used in the context of T_C - k -sets, we show a slightly stronger bound for ranges (induced by translations) that *do not contain points on their boundaries*.

Lemma 6.5.2 (lower bound on number of ranges, probabilistic). *Let $C \subseteq \mathbb{R}^2$ be the interior of a convex body. Let $A \subseteq \mathbb{R}^2$ be a compact set such that $2C \subseteq A$.¹ Let X be a set of n iid uniformly random points in A . Let $t > 0$. Then there exists a constant $c_{6.5.2} > 0$ that depends only on A , C and t such that with probability at least $1 - n^{-t}$,*

$$|\{X \cap (x + C) : x + C \subseteq A, X \cap \text{bd}(x + C) = \emptyset\}| \geq c_{6.5.2} \left(\frac{n}{\log n} \right)^2.$$

PROOF. Let $X = \{X_1, \dots, X_n\}$. Let B be the ball with center 0 and radius 1. Without loss of generality (up to scaling and translation), $B \subseteq C$.

We will first construct a packing of $n^2/(c \log n)^2$ translations of C with centers in B with area of pairwise symmetric difference at least about $\log(n)/n$, for some $c > 0$ to be determined later. Let G be an $n/(c \log n)$ -by- $n/(c \log n)$ grid of points with gap $(c \log n)/n$ between adjacent rows and columns of points and contained in B . Every pair of points in G is then at distance at least $(c \log n)/n$, and therefore for all $x, y \in G$ with $0 < \|x - y\| \leq 1$ we have $\text{area}((x + C)\Delta(y + C)) =$

¹The assumption that $2C \subseteq A$ guarantees that the translations of C from our grid G in the proof are contained in A so that the probability computation goes through.

$2\left(\text{area}(C) - \text{area}(C \cap ((y-x) + C))\right) \geq \text{area}(C)\|y-x\| \geq \text{area}(C)(c \log n)/n$ (using Lemma 6.5.1).
The bound extends to all $x, y \in G$ with $x \neq y$ by monotonicity.

We will now show that with probability at least $1 - o(1)$ each $x + C$ with $x \in G$ induces a different range on X . It is enough to show that for all $x, y \in G$ with $x \neq y$ we have $((x+C)\Delta(y+C)) \cap X \neq \emptyset$. Setting $c = (t+2) \text{area}(A)/\text{area}(C)$, the probability of this event for some x, y is

$$\begin{aligned} \mathbb{P}((\forall i \in [n]) X_i \notin (x+C)\Delta(y+C)) &\leq \left(1 - \frac{\text{area}((x+C)\Delta(y+C))}{\text{area}(A)}\right)^n \\ &\leq \left(1 - \frac{\text{area}(C)}{\text{area}(A)} \frac{(c \log n)}{n}\right)^n \\ &\leq e^{-(t+2) \log n} \\ &= 1/n^{t+2}. \end{aligned}$$

Thus, with probability at least $1 - n^2/n^{t+2} = 1 - 1/n^t$ our event holds for all pairs x, y .

Finally, $X \cap \cup_{x \in G} \text{bd}(x+C) = \emptyset$ a.s. The claim follows. \square

We now state and prove our probabilistic lower bound for T_C - k -sets for some k proportional to n :

Theorem 6.5.3. *Let $C \subseteq \mathbb{R}^2$ be the interior of a convex body. Let $A \subseteq \mathbb{R}^2$ be a compact set such that $2C \subseteq A$. Let X be a set of n iid uniformly random points in A . Let*

$$a'_k(X) := |\{X \cap (x+C) : x+C \subseteq A, |X \cap (x+C)| = k, X \cap \text{bd}(x+C) = \emptyset\}|$$

(that is, $a'_k(X)$ is the number of T_C - k -sets of X induced by translations of C contained in A). Let $p = \text{area}(C)/\text{area}(A)$. Then there exists a function $k(n)$ such that $\mathbb{E}(a_{k(n)}(X)) \geq \mathbb{E}(a'_{k(n)}(X)) \geq \Omega(n^{3/2}/(\log n)^{5/2})$ and $|k(n) - pn| \leq O(\sqrt{n \log n})$ (where function $k(n)$ and the constants in O, Ω depend only on A and C).

PROOF. Let (A, \mathcal{R}) be the set system where \mathcal{R} is the family of translations of C contained in A . From Theorem 6.4.3 we have that the growth function $s(n)$ of (A, \mathcal{R}) satisfies $s(n) \leq n^3$.

Fix n . For $S \subseteq A$, let $\hat{P}(S) = |X \cap S|/n$. From Theorem 2.4.6 (VC's uniform convergence theorem)² we have, for $n \geq 2/\epsilon^2$,

$$(6.1) \quad \mathbb{P}\left(\sup_{S \in \mathcal{R}} |\hat{P}(S) - p| > \epsilon\right) \leq 4s(2n)e^{-\epsilon^2 n/8}.$$

Set $\epsilon = 4\sqrt{\frac{3 \log 2n}{n}}$ so that the rhs is at most $1/n^3$. If we denote by G the complement of the event in (6.1), we have $\mathbb{P}(G) = 1 - o(1)$.

Let $\mathcal{R}_X = \{X \cap S : S \in \mathcal{R}, X \cap \text{bd } S = \emptyset\}$. From Lemma 6.5.2 with $t = 1$ and notation $f(n) = c_{6.5.2}(n/\log n)^2$, we get

$$(6.2) \quad \mathbb{P}(|\mathcal{R}_X| \geq f(n)) \geq 1 - o(1).$$

Let H denote the event in (6.2).

To conclude, we will show that there is a value $k(n)$ that is independent of X and makes $\mathbb{E}(a'_k(X))$ large. We have $\mathbb{P}(X \in G \cap H) = 1 - o(1)$. Also for $X \in G \cap H$ we have

$$\sum_{k \in [pn - n\epsilon, pn + n\epsilon]} a'_k(X) = |\mathcal{R}_X| \geq f(n).$$

Therefore

$$\mathbb{E}\left(\sum_{k \in [pn - n\epsilon, pn + n\epsilon]} a'_k(X)\right) \geq f(n) \mathbb{P}(X \in G \cap H).$$

Reordering, $\sum_{k \in [pn - n\epsilon, pn + n\epsilon]} \mathbb{E}(a'_k(X)) \geq f(n) \mathbb{P}(X \in G \cap H)$. Thus, there exists $k(n) \in [pn - n\epsilon, pn + n\epsilon]$ such that $\mathbb{E}(a'_k(X)) \geq f(n) \mathbb{P}(X \in G \cap H)/(2n\epsilon + 1)$. That is (using $n\epsilon = O(\sqrt{n \log n})$), $\mathbb{E}(a'_k(X)) \geq \Omega(n^{3/2}/(\log n)^{5/2})$.

For the bound on $\mathbb{E}(a_k(X))$, from the definitions we have $a_k(X) \geq a'_k(X)$. □

²In order to apply VC's uniform convergence theorem, we need to verify that the function $\sup_{S \in \mathcal{R}} (\cdot)$ as defined in Eq. (6.1) is measurable, i.e., that it is a random variable. This can be verified by observing that \mathcal{R} is a *permissible* class of subsets of A . See [Pol84, Appendix C] for the definition of permissible classes and a proof of the measurability of suprema in this context. One can see that the class \mathcal{R} is permissible by indexing it by translation and verifying that the requirements for permissibility are met.

CHAPTER 7

Open questions

7.1. The real degree of algebraic varieties

The degree of complex algebraic varieties is an extremely well-studied and useful definition. Of course, this definition of degree also applies to real algebraic varieties. The problem is that the degree does not always capture the geometry of a given real variety. For example, the degree of the variety $V(x^4 - y)$ is 4. But if you just look at the real points, you would say that the degree should be 2. Here we suggest a new definition of degree for real algebraic varieties (which depends only on the real points) and suggest how it may be related to the property of a real variety being weakly k -neighborly. We remark that a notion of degree for real varieties was suggested in [NR09]. However, the notion of degree given there does not line up with the definition we suggest here.

Definition 7.1.1. *Let $V \subset \mathbb{R}^p$ be a d -dimensional real algebraic variety which contains smooth real points. The real degree of V is the maximum integer D so that there exists an open subset O of the set of all $(p - d)$ -flats in \mathbb{R}^p with the property that for every flat L in O , $|L \cap V(\mathbb{R})| = D$.*

We say that a d -dimensional variety $V \subset \mathbb{R}^p$ is a variety of *minimal real degree* if the real degree of V is $p - d + 1$. Observe that the parametric real algebraic curves of minimal real degree are precisely the curves (called p -order curves or curves of order p) which are studied by Sturmfels in [Stu87] because of their connection to neighborly polytopes.

We need to make one more definition which is inspired by Section 6.5:

Definition 7.1.2. *A d -dimensional algebraic variety $V \subset \mathbb{R}^p$ is maximally weakly neighborly if V is weakly k -neighborly and $p = 2k + d - 1$.*

It is a result of Cordovil and Duchet that a (parameterized) real algebraic curve is maximally weakly neighborly if and only if it is a curve of minimal real degree [CD00, Proposition 3.6]. This equivalence might hold for higher dimensional varieties as well. That is, it might be true that

a real variety (possibly with some assumptions on the variety) is maximally weakly neighborly if and only if it is a variety of minimal real degree. The same equivalence may also be true with real varieties replaced by parametric surfaces and hypersurfaces if one extends the definitions of maximally weakly neighborly and minimal real degree to arbitrary parametric (hyper)surfaces in \mathbb{R}^d .

7.2. Conjecture on generally k -neighborly embeddings/manifolds

Conjecture 4.2.5 remains open. Of course resolving the conjecture in full generality would be ideal, but it may be worthwhile to focus on other special cases instead. For example, one might be able to prove the conjecture for analytic manifolds or for smooth manifolds using tangency properties.

7.3. Lower bounds for the algebraic k -set problem

In analogy with the original k -set problem we ask: Is there a polynomial map of the plane into some higher dimensional space which induces a natural set system $(\mathbb{R}^2, \mathcal{F})$ for which the maximum number of \mathcal{F} - k -sets for a set of n points is $n^\ell e^{\Omega(\sqrt{\log n})}$ and $o(n^{\ell+1})$ for some integer $\ell \geq 2$? Or just $\Omega(n^\ell \log n)$ and $o(n^{\ell+1})$? A candidate that we do not fully understand in this context is the map $(x, y) \mapsto (x, y, xy)$ (or, equivalently up to a linear transformation, $(x, y) \mapsto (x, y, x^2 - y^2)$).

7.4. Bound on the number of points of intersection of $Z(f)$ and the k -edge graph

Can the answer to Question 5.3.1 given in Lemma 5.3.8 be improved? We believe it may be possible to improve the bound to $O(nr)$ for the following reason: The k -edge graph G_k of a set of n points behaves somewhat like a degree n algebraic curve when it comes to intersecting it with a line. In particular, a degree n algebraic curve and the k -edge graph of a set of n points both have the property that a line can intersect them at most n times unless the line intersects them infinitely many times. One might expect this phenomenon to also hold for arbitrary algebraic curves, not just lines. If this is the case, the bound on the number of intersection points between a degree r algebraic curve and the k -edge graph of a set of n points should be $O(nr)$ as in the case of the intersection of a degree r algebraic curve and a degree n algebraic curve.

7.5. Polynomial partitioning for k -sets in higher dimensions

The polynomial partitioning theorem becomes more powerful in higher dimensions. It may be possible to apply it to the k -set problem in dimensions higher than 2. The issue is that there is no analogue of the convex chains decomposition in dimension higher than 2, so this would likely require the discovery of a new property of k -sets or k -facets.

Bibliography

- [AACS98] P. K. Agarwal, B. Aronov, T. M. Chan, and M. Sharir, *On levels in arrangements of lines, segments, planes, and triangles*, Discrete Comput. Geom. **19** (1998), no. 3, Special Issue, 315–331, Dedicated to the memory of Paul Erdős.
- [ABFK92] N. Alon, I. Bárány, Z. Füredi, and D. J. Kleitman, *Point selections and weak ϵ -nets for convex hulls*, Combin. Probab. Comput. **1** (1992), no. 3, 189–200.
- [Ard04] F. Ardila, *The number of halving circles*, The American Mathematical Monthly **111** (2004), no. 7, 586–591.
- [BCR98] J. Bochnak, M. Coste, and M.-F. Roy, *Real algebraic geometry*, Springer, 1998.
- [BHP08] S. Buzaglo, R. Holzman, and R. Pinchasi, *On s -intersecting curves and related problems*, Proceedings of the 24th ACM Symposium on Computational Geometry, College Park, MD, USA, June 9–11, 2008 (M. Teillaud, ed.), ACM, 2008, pp. 79–84.
- [BS94] I. Bárány and W. Steiger, *On the expected number of k -sets*, Discrete Comput. Geom. **11** (1994), no. 3, 243–263.
- [CD00] R. Cordovil and P. Duchet, *Cyclic polytopes and oriented matroids*, Eur. J. Comb. **21** (2000), no. 1, 49–64.
- [CFSS14] N. Chevallier, A. Fruchard, D. Schmitt, and J. Spehner, *Separation by convex pseudo-circles*, 30th Annual Symposium on Computational Geometry, SOCG’14, Kyoto, Japan, June 08 - 11, 2014 (S. Cheng and O. Devillers, eds.), ACM, 2014, p. 444.
- [CFSS20] ———, *Separation by convex pseudo-circles*, Discrete & Computational Geometry (2020).
- [Cla87] K. L. Clarkson, *New applications of random sampling in computational geometry*, Discrete & Computational Geometry **2** (1987), 195–222.
- [Cla04] K. L. Clarkson, *On the expected number of k -sets of coordinate-wise independent points*, https://kenclarkson.org/cwi_ksets/p.pdf, 2004, Manuscript.
- [CP86] B. Chazelle and F. P. Preparata, *Halfspace range search: An algorithmic application of k -sets*, Discret. Comput. Geom. **1** (1986), 83–93.
- [CS89] K. L. Clarkson and P. W. Shor, *Application of random sampling in computational geometry, II*, Discrete & Computational Geometry **4** (1989), 387–421.
- [CSY87] R. Cole, M. Sharir, and C.-K. Yap, *On k -hulls and related problems*, SIAM J. Comput. **16** (1987), no. 1, 61–77.

- [Der82] Der-Tsai Lee, *On k -nearest neighbor Voronoi diagrams in the plane*, IEEE Transactions on Computers **C-31** (1982), no. 6, 478–487.
- [Dey97] T. K. Dey, *Improved bounds on planar k -sets and k -levels*, Proceedings 38th Annual Symposium on Foundations of Computer Science (1997), 156–161.
- [Dey98] ———, *Improved bounds for planar k -sets and related problems*, Discrete & Computational Geometry **19** (1998), no. 3, 373–382.
- [ELSS73] P. Erdős, L. Lovász, A. Simmons, and E. G. Straus, *Dissection graphs of planar point sets*, A survey of combinatorial theory (Proc. Internat. Sympos., Colorado State Univ., Fort Collins, Colo., 1971), 1973, pp. 139–149.
- [GK15] L. Guth and N. H. Katz, *On the Erdős distinct distances problem in the plane*, Ann. of Math. (2) **181** (2015), no. 1, 155–190.
- [GP10] V. Guillemin and A. Pollack, *Differential topology*, AMS Chelsea Publishing, Providence, RI, 2010, Reprint of the 1974 original.
- [Gru07] P. M. Gruber, *Convex and discrete geometry*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 336, Springer, Berlin, 2007.
- [Gut16a] L. Guth, *Polynomial methods in combinatorics*, University Lecture Series, vol. 64, American Mathematical Society, Providence, RI, 2016.
- [Gut16b] ———, *A restriction estimate using polynomial partitioning*, J. Amer. Math. Soc. **29** (2016), no. 2, 371–413.
- [Har95] J. Harris, *Algebraic geometry*, Graduate Texts in Mathematics, vol. 133, Springer-Verlag, New York, 1995, A first course, Corrected reprint of the 1992 original.
- [HP11] S. Har-Peled, *Geometric approximation algorithms*, Mathematical Surveys and Monographs, vol. 173, American Mathematical Society, Providence, RI, 2011.
- [Hum40] P. M. Hummel, *A note on Stirling’s formula*, Amer. Math. Monthly **47** (1940), no. 2, 97–99.
- [Kir92] F. Kirwan, *Complex algebraic curves*, London Mathematical Society Student Texts, vol. 23, Cambridge University Press, Cambridge, 1992.
- [KMS12] H. Kaplan, J. Matoušek, and M. Sharir, *Simple proofs of classical theorems in discrete geometry via the Guth-Katz polynomial partitioning technique*, Discrete Comput. Geom. **48** (2012), no. 3, 499–517.
- [KW08] G. Kalai and A. Wigderson, *Neighborly embedded manifolds*, Discrete & Computational Geometry **40** (2008), no. 3, 319–324.
- [Lov71] L. Lovász, *On the number of halving lines*, Ann. Univ. Sci. Budapest, Etovos, Sect. Math. **14** (1971), 107–108.
- [LR20] B. Leroux and L. Rademacher, *Algebraic k -sets and generally neighborly embeddings*, 2020, arXiv:1912.03875.

- [LR21] ———, *Improved bounds for the expected number of k -sets*, 2021, arXiv:2106.04782.
- [Mat02] J. Matoušek, *Lectures on discrete geometry*, Graduate texts in mathematics, vol. 212, Springer, 2002.
- [Mil64] J. Milnor, *On the Betti numbers of real varieties*, Proc. Amer. Math. Soc. **15** (1964), 275–280.
- [MSRB71] P. McMullen, G. C. Shephard, J. E. Reeve, and A. Ball, *Convex polytopes*, vol. 3, CUP Archive, 1971.
- [MSSW06] J. Matoušek, M. Sharir, S. Smorodinsky, and U. Wagner, *k -sets in four dimensions*, Discrete Comput. Geom. **35** (2006), no. 2, 177–191.
- [NR09] J. Nie and K. Ranestad, *Algebraic degree of polynomial optimization*, SIAM J. Optim. **20** (2009), no. 1, 485–502.
- [NT10] M. Naszódi and S. Taschuk, *On the transversal number and vc -dimension of families of positive homothets of a convex body*, Discrete Mathematics **310** (2010), no. 1, 77–82.
- [Pol84] D. Pollard, *Convergence of stochastic processes*, Springer Series in Statistics, Springer-Verlag, New York, 1984.
- [Rut00] J. W. Rutter, *Geometry of curves*, Chapman & Hall/CRC Mathematics, Chapman & Hall/CRC, Boca Raton, FL, 2000.
- [Sch14] R. Schneider, *Convex bodies: the Brunn-Minkowski theory*, second expanded ed., Encyclopedia of Mathematics and its Applications, vol. 151, Cambridge University Press, Cambridge, 2014.
- [She20] A. Sheffer, *Polynomial methods and incidence theory*, <http://faculty.baruch.cuny.edu/ASheffer/000book.pdf>, 2020, Accessed: 2021-02-05.
- [SST01a] M. Sharir, S. Smorodinsky, and G. Tardos, *An improved bound for k -sets in three dimensions*, Discrete Comput. Geom., vol. 26, Springer, 2001, pp. 195–204, ACM Symposium on Computational Geometry (Hong Kong, 2000).
- [SST01b] M. Sharir, S. Smorodinsky, and G. Tardos, *An improved bound for k -sets in three dimensions*, Discret. Comput. Geom. **26** (2001), no. 2, 195–204.
- [ST42] A. H. Stone and J. W. Tukey, *Generalized “sandwich” theorems*, Duke Math. J. **9** (1942), 356–359.
- [STC04] J. Shawe-Taylor and N. Cristianini, *Kernel methods for pattern analysis*, Cambridge University Press, New York, NY, USA, 2004.
- [Stu87] B. Sturmfels, *Cyclic polytopes and d -order curves*, Geometriae Dedicata **24** (1987), no. 1, 103–107.
- [Tót01] G. Tóth, *Point sets with many k -sets*, Discrete & Computational Geometry **26** (2001), no. 2, 187–194.
- [VC71] V. Vapnik and A. Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*, Theory of Probability and its Applications **16** (1971), 264 – 280.
- [Wag08] U. Wagner, *k -sets and k -facets*, Surveys on discrete and computational geometry, Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 443–513.
- [Whi57] H. Whitney, *Elementary structure of real algebraic varieties*, Ann. of Math. (2) (1957), 545–556.
- [Zie95] G. M. Ziegler, *Lectures on polytopes*, Springer-Verlag, New York, 1995.