1 A Happy Ending

In the early 1930s, Hungarian mathematician Esther Klein made a discovery that, despite its apparent simplicity, would kick off two major lines of research in mathematics. Klein observed that every set of five points in the plane has either three points in a line or four points in a convex quadrilateral. This became one of the first results in the two fields of discrete geometry (the study of combinatorial properties of geometric objects such as points in the Euclidean plane, and the subject of this book) and Ramsey theory (the study of the phenomenon that unstructured mathematical systems often contain highly structured subsystems).

Klein's observation can be proven by a simple case analysis that considers how many of the points belong to their *convex hull*. The convex hull is a convex polygon, having some of the given points as its vertices and containing the others. It can be defined mathematically in many ways, for instance as the smallest-area convex polygon that contains all of the given points or as the largest-area simple polygon whose vertices all belong to the given points. The convex hull of points that are not all on a line always has at least three vertices (for otherwise it could not enclose a nonzero area) and, for five given points, at most five vertices. If it has five vertices, any four of them form a convex quadrilateral, and if it has four vertices then it is a convex quadrilateral. The remaining possibility for the convex hull is a triangle, with the other two points either part of a line of three points or inside the triangle. When both points are inside, and the line through them misses the triangle vertices, it also misses one side of the triangle. In this case the two interior points and the two points on the missed side form a convex quadrilateral (Figure 1.1).

The challenge of extending and generalizing this observation was taken up by two of Klein's friends, Paul Erdős and George Szekeres. They proved that,

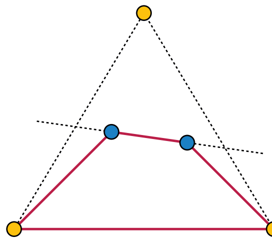

Figure 1.1. Five points with no three in line always contain a convex quadrilateral, either four vertices of the convex hull or (as shown here) two vertices of a triangular convex hull and two interior vertices. When the convex hull is a triangle, the line through the remaining two points has two triangle vertices on one of its sides, and those four points are in convex position.

for every k, a convex k-gon can be found in all large enough sets of points, as long as no three of the points lie on a line.[1] Klein later married Szekeres, and their marriage is commemorated in the name of Erdős and Szekeres's result: the happy ending theorem.

Erdős and Szekeres published two proofs of their theorem, one of which showed that every 4^k points in general position (meaning, no three in a line) contain a convex k-gon.[2] However, this gives only a loose estimate. For instance, it would tell us that we need $4^4 = 256$ points to guarantee the existence of a convex quadrilateral, many more than the five points of Klein's observation. Therefore, it became of interest to determine more precisely how many points are needed to ensure the existence of a convex k-gon.

In their original work on this problem, Erdős and Szekeres conjectured that many fewer points, $2^{k-2} + 1$ of them, would already force a convex k-gon to exist. Later, they constructed sets of 2^{k-2} points with no three in line and no convex k-gon, so if true their conjecture would be as tight as possible.[3] For example, some sets of eight points have no convex pentagon, matching the formula as $8 = 2^{5-2}$ (Figure 1.2). We detail their construction in Section 11.1. Tightening the gap between this construction and the 4^k upper bound remained open until in a recent breakthrough Suk (2017) proved that for sufficiently large k, every $2^{k+6k^{2/3}\log k}$ points in general position contain a convex k-gon. Although this does not settle the conjecture of Erdős and Szekeres, it has the correct leading term in the exponent and brings the upper and lower bounds much closer.[4]

1 Erdős and Szekeres (1935).
2 More precisely, they showed that every set of $\binom{2k-4}{k-2} + 1 \leq 4^k$ points has a convex k-gon. Here $\binom{2k-4}{k-2}$ is a *binomial coefficient*, the number of ways of choosing $k-2$ elements from a set of $2k-4$ elements.
3 Erdős and Szekeres (1960).
4 Suk writes that Gábor Tardos has further improved the low-order term in the exponent of this bound. For an intuitive overview of Suk's proof, see Hartnett (2017).

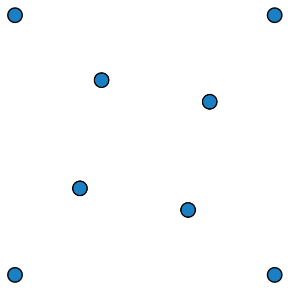

Figure 1.2. Eight points in general position that do not contain the vertices of a convex pentagon.

Open Problem 1.1 (the happy ending problem)

Does there exist an integer k, and a set of $2^{k-2}+1$ points in the plane with no three in a line and no k forming a convex k-gon?

Two features of the happy ending theorem, and of the convex subsets of point sets that it describes, are of particular interest to us. First, the size of the largest convex subset of a set of points is *monotone*: if you remove points from the set, then its largest convex subset can only decrease in size or stay the same, but it can never grow. Second, the convex subsets of point sets are insensitive to the precise locations of the points. If you move the points around the plane in a continuous motion, being careful only to never let three of them line up, you cannot create new convex polygons nor destroy the ones that are already present. Another way to express the same insensitivity is that the convex subsets of a point set depend only on the *orientations* of the points: which triples are in clockwise order, which counterclockwise, and which collinear. Many other problems in discrete geometry share these characteristics: they involve monotone properties of finite point sets that depend only on the orientations of the points. Problems of this type are the subject of this book.

Cambridge University Press
978-1-108-42391-5 — Forbidden Configurations in Discrete Geometry
David Eppstein
Excerpt
More Information

2 Overview

Many algorithmic and combinatorial problems concerning finite sets of points have been studied in discrete and computational geometry. Often, the answers to these problems depend only on knowing, for each three points, whether they are in clockwise order, counterclockwise order, or lie on a single line. It is safe, for these problems, to throw away the coordinates of the points and retain only their *configuration*, which tells us this ordering information for each triple of points. In many cases, in addition, the property or quantity to be studied behaves predictably under the removal of points. If removing a point can never cause a quantity of interest to increase, we call that quantity *monotonic*. Our goal in this work is to provide a systematic study of the monotonic properties of configurations.

Several old and colorfully named puzzles and games fit this pattern:

- The happy ending theorem was famously given its name after its proof led to the marriage of two of the mathematicians who discovered it, Esther Klein and George Szekeres. It is about how many points are needed (no three in a line) to ensure the existence of a convex polygon with a given number of corners. We described it already in Chapter 1.
- The orchard-planting problem, which we describe in Section 8.1, dates back to the early nineteenth century. It asks how many rows of three trees one can form by planting an orchard with a given number of trees.
- We describe the no-three-in-line problem in Section 9.1. It was first posed in terms of placing 16 pawns on a chessboard, so that no three of them line up with each other. For this problem, it is important to consider all directions of lines, not just the horizontal, vertical, and diagonal directions of the chessboard.

Other topics of more serious past research also concern monotonic properties of configurations. They include searching for triples of points that all lie on a single line (Chapter 8), grouping points into clusters of collinear points (also Chapter 8), finding convex polygons within sets of points (Chapter 11), partitioning sets of points into nested convex polygons ("onion layers," Section 12.4), estimating the center of a cloud of points in a way that is insensitive to perturbations of the points (Section 12.7), perturbing sets of points so their distances or coordinates are all integers (Chapter 13), using sets of points as vertices to draw planar graphs (Chapter 16), and finding paths that no line can cross many times (Section 17.1).

Our study of the monotone properties and parameters of configurations looks at them from the following points of view.

Forbidden configurations. Each monotone property, and each value of a monotone parameter, can be characterized by its *forbidden configurations* or *obstacles*. These are the configurations that do not have the property (or that have too large a value) but for which all subconfigurations do have it. The properties of any particular configuration can be read off from whether it contains any of these obstacles.

For the properties and parameters we consider, we ask: are there a finite number of forbidden configurations? If so, can we describe them all, or bound their size as a function of k?

Computational complexity. Can the given property or parameter be computed in polynomial time, or is it NP-hard? If a parameter is hard, how well can it be approximated, and how efficiently can we compute small parameter values? Can we distinguish point sets with a property from sets far from having the property by examining only small samples of the set, or is it necessary to test the whole input?

Both the problems of computing small parameter values and of using samples to test properties are closely related to the existence of finitely many obstacles. As we will see, algorithms using a technique called *kernelization* can often be used to bound the size of the obstacles for a parameter. And the size of a sample needed to test any property can be bounded in terms of the size of the obstacles for the property. However, we will also see that in some cases a parameter may have a constant number of obstacles for each parameter value but still be hard to compute, even for bounded parameter values.[1]

Inequalities. How are different monotone parameters related to each other? Which ones are bounded above or below by functions of each other?

Well-quasi-ordering. For which families of configurations do all monotone properties on that family have a finite number of obstacles? The families

[1] For the connection to sampling, see Section 6.5. For a parameter that is hard for bounded parameter values, see Theorem 7.7.

for which this is true are called *well-quasi-ordered*. The family of all configurations is not well-quasi-ordered, but some of its subfamilies are, and we investigate which ones.

Combinatorial enumeration. How many configurations have a given property? The number of all n-point configurations is known to grow exponentially in $n \log n$.[2] Is the growth rate of the configurations with a given property or obstacle significantly smaller?

Along the way we will also collect a large menagerie of monotone properties and monotone parameters of configurations. In addition we investigate the complexity of testing whether one configuration is part of another larger one. This can be done quickly when the size of the smaller configuration is fixed, but is much harder when its size can vary.

Although much of this work should be readable by nonspecialists, some of it remains technical. Most chapters place the more generally accessible aspects of their subject in the earlier parts of the chapter, and the more technical aspects later, so readers who find some material difficult should feel free to skip ahead to the start of the next chapter where it will likely be easier going again. It is generally safe to skip past proofs, at least on a first reading, and many of the results in the earlier chapters of the book have proofs that we have delayed until the later chapters. Chapters 3–5 give the main definitions that we use, of configurations, subconfigurations, monotone parameters, and monotone properties, and some important general results based on these definitions. Chapters 6 and 7 introduce the algorithmic study of configurations. The remaining chapters are largely independent from each other. They discuss different subtopics of discrete geometry that all involve monotone parameters of configurations.

Much of this work has analogies with the theory of graphs and subgraphs (or graphs and minors) and with the theory of permutations and permutation patterns. For instance, the questions raised above about combinatorial enumeration are analogous to the Stanley–Wilf conjecture for permutations, proved by Marcus and Tardos (2004), according to which forbidding a single permutation pattern reduces the number of permutations in a permutation class from factorial to single exponential. We will exploit the analogy to permutations and permutation patterns in Chapter 14 by finding general methods of translating permutations into configurations. Similarly in Chapter 15 we translate graphs into configurations. We use these translations both positively, to develop efficient algorithms for configurations, and negatively, to show that certain problems on configurations are hard to compute.

We conclude with Chapter 18, which summarizes the results from these chapters and provides a road map of the relations between parameters from different chapters.

[2] Goodman and Pollack (1986).

Because of our focus on monotone properties of configurations, there are many aspects of discrete geometry that are beyond the scope of our work. For instance, there have been recent breakthroughs on the number of distinct distances that any point set in the plane must have,[3] and on related problems of counting incidences of geometric objects, that we do not cover here.

[3] Guth and Katz (2015).

3 Configurations

We begin our study with an examination of the different ways we can place small numbers of points in the plane, and what it means for two sets of points to be different.

3.1 Small Configurations

In how many different ways can we place n points in the plane? With a list of all of the possible placements, we could prove statements such as Klein's observation about convex quadrilaterals in five-point sets, automatically, merely by checking all the cases. When n is small enough, we can provide an explicit answer.

Example 3.1

There are two different ways of placing three points in the plane: they may either lie on a line, or they may form a triangle.

Four points may be arranged in four different ways:

- a four-point line
- three points on a line and one off the line
- a triangle containing one point, or
- a convex quadrilateral.

Figure 3.1 depicts these three-point and four-point configurations.

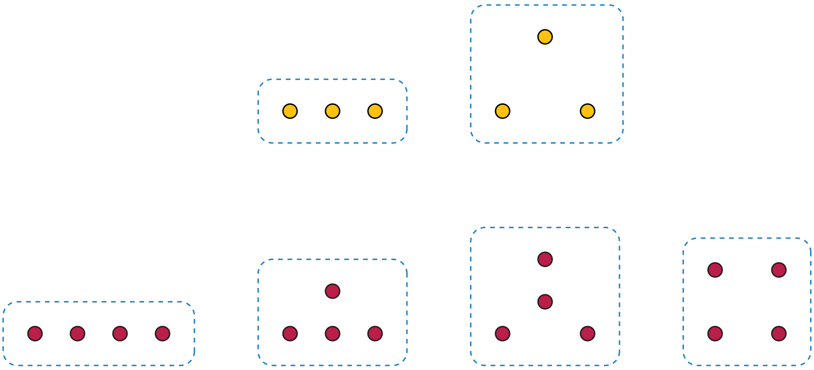

Figure 3.1. Two ways of arranging three points (yellow, top) and four ways of arranging four points configurations (red, bottom).

So far, we have used only an intuitive notion for what it means for two sets of points to be the same or different. But when we get to five points, we already need to define more precisely what we mean. Figure 3.2 depicts 13 sets of five points. Are they all different from each other? Two of these are mirror images: should that count as two different sets of points or as two different views of a single way of placing five points?

We will give a more precise definition of what it means for sets of points to be the same or different in the next section. For the definition we use, mirror images are not (in general) considered to be the same, so there are indeed 13

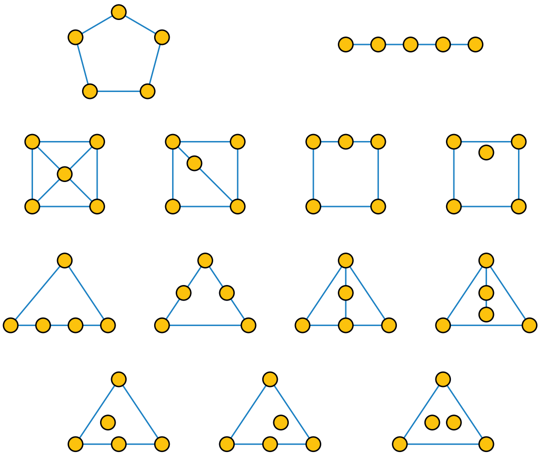

Figure 3.2. The 13 possible five-point configurations. Each configuration is shown together with its convex hull edges and with any line segments that pass through three or more collinear points.

different ways of arranging five points. Open Problem 3.5, later in this chapter, formalizes the problem of counting configurations of different sizes.

The numbers of sets of n points grow very quickly, as $c^{n \log n}$ for a constant $c > 1$.[1] Aichholzer et al. (2002) used computer searches to establish a database of small sets of points.[2] To help control the size of the database, they exclude sets that have three points on a line and count mirror-image sets of points as equivalent. Despite these restrictions, the sets in their database still grow as $c^{n \log n}$, but with a smaller c. Their numbers are:

n:	1	2	3	4	5	6	7	8	9	10	…
#:	1	1	1	2	3	16	135	3315	158817	14309547	…

For instance the two configurations they count with $n = 4$ are the two on the lower right of Figure 3.1.

3.2 Orientations and Order Types

To test whether two sets of points are placed in the same way or differently, we use *order types*. Intuitively, the order type of a set of points in the plane describes how the points are positioned with respect to each other: which points are on lines with each other, and how do the lines through pairs of points split the remaining points? We define these concepts more formally below. Our definitions will allow us to define the shape of a set of points in an abstract way that does not depend on how the points are scaled or rotated or on other inessential properties.

The building blocks of order types are the orientations of triples of points.[3]

Definition 3.2

We define the *orientation* of an ordered triple of points in the Euclidean plane to be one of the three numbers $+1$, -1, or 0. It is $+1$ if the three points form the vertices of a nonzero-area triangle, listed by the triple in clockwise order around the triangle. It is -1 if they are in counterclockwise order, or 0 if they are collinear.

Two cyclic permutations of the same triple have the same orientation, and reversing a triple causes its orientation to be negated. Therefore, if we know the

[1] Goodman and Pollack (1986).
[2] This database is online at www.ist.tugraz.at/aichholzer/research/rp/triangulations/ordertypes/.
[3] It is not possible to break down this information further, into locally defined structures on pairs of points; see Balko et al. (2017).