Introduction

> "When you have answered the question, it's time to question the answer."
>
> Paul Fjelstad

Good problems invite exploration and play, while needing perseverance. They also engender new problems and even new areas of mathematics. Discrete geometry developed in the twentieth century from problems that, while intriguing, seemed on the edge of traditional geometry. This book uses variations of a number of problems to lead to a deeper understanding of this relatively new area.

Discrete geometry studies arrangements of different numbers of points, lines, and other familiar objects, often looking for optimal arrangements or counting the number of ways of making these arrangements. Problems include counting distances determined by a set of points or placing "guard points" to "see" all other points in a given region. Others challenge us to divide a polygon into triangular regions or pack circles efficiently among many other problems. Many discrete geometry questions start from pure mathematical curiosity. More recently many of them have found application in computer imaging and other areas. Because of these applications, discrete geometry also involves algorithms to solve various aspects of these problems.

The answers to initial questions in this area often spark a variety of related questions. We fully embrace this idea of generating new questions as a unifying theme of this book, embodied in the Fjelstad quote above that starts our explorations. The following descriptions of the book's chapters illustrate this process.

- Chapter 1 provides a first layer of easily posed questions, and challenges you to play with them and perhaps even solve them.

- Chapter 2 answers the questions from Chapter 1, introducing some mathematical ideas along the way. Even more, it takes to heart the Fjelstad quote starting this introduction by posing new problems that are variations of the problems of Chapter 1.

- Chapter 3 answers the problems of Chapter 2, provides related mathematical ideas and, continuing the theme of the book, questions those answers with more variations.

- Chapter 4 brings closure by answering the variations of Chapter 3 and indicating a broader view, including what is known about some of the topics in this area.

- The material after Chapter 4 starts with answers to the exercises not answered earlier. After that are suggestions of books and articles at an accessible level for further exploration, together with a short description of each one. (Of course, you can find much more information on these topics on many sites on the internet. However, determining whether a given site is a good source is much harder.) We use the mathematical form of referencing in the text, such as [6, 21], which refers to page 21 of the sixth item (the book *Excursions into Mathematics*) in the Bibliography section, which appears after the Suggested Readings. The Index rounds out the end material.

This book's problems seek to challenge anyone excited about mathematics and doesn't assume any background beyond high school mathematics. The explanations also require no more background, although the topics range far beyond the high school curriculum. We assume a familiarity with common geometrical and algebraic concepts, rather than burdening the text with a more formal presentation. We will define less-common terms, which appear in italics in their definition. The reasoning will vary from computations and informal arguments to more careful, but still accessible, proofs to fit each context.

Discrete geometry can improve valuable visualization skills because it asks different kinds of questions from those asked in traditional geometry classes. The problems in discrete geometry, like any area of mathematics, will also improve your problem solving skills, a valuable asset for anybody. The process of "questioning the answer" and the variations it spawns raises problem solving to a new level. From a mathematician's point of view, the variations provide more important benefits: they help deepen mathematical reasoning and geometrical intuition, and they introduce one way mathematicians find new research questions. Indeed, some variations of problems we consider lead to open research questions or recent results. Mathematical research often develops from seeing patterns and making conjectures about why those patterns happen. I hope you will grow mathematically by wrestling with the variations in this book, seeing patterns in them, making conjectures about them, solving them, and even making up your own variations to explore! But the main motivation for writing this book is that I found these problems and variations fun to play with—and I hope you find much pleasure in pondering them as well.

Acknowledgments

My thanks to the Anneli Lax New Mathematical Library Editorial Board of the Mathematical Association of America and especially the chair, Bruce Torrence, for many helpful suggestions and guidance. I greatly appreciate the expertise and professionalism of my editor, Sergei Gelfand, and his team from the American Mathematical Society. My wife, Jennifer Galovich, has given the loving support and sympathetic ear that I needed and valued throughout the writing. I remember with deep fondness the inspiration of my college mentor, Paul Fjelstad, whose creativity and enthusiasm guided me throughout my career. This book is dedicated in his memory.

1

Beginning Explorations

Discrete geometry, like all of mathematics, has developed piecemeal. It is still new enough to have unsolved problems that seem just a few steps from introductory ones. The problems of this chapter seek to whet your appetite. Think about them—draw pictures for small cases, look for patterns, make conjectures, and try to solve them. Complete answers are great, but so are partial answers, guesses, and even mistakes. Recall: mathematics is not a spectator sport; so enjoy engaging with these problems. We'll revisit these problems in later chapters, providing answers, relevant mathematics, and especially new questions probing these areas more deeply.

1.1 Lines and Regions

A line drawn on a sheet of paper splits the paper into two regions. Two lines determine either three or four regions, as indicated in Figure 1.1.

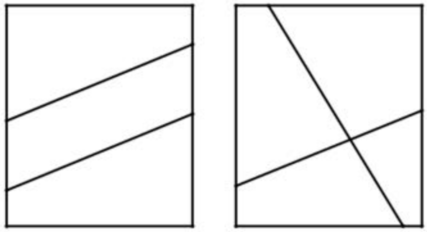

Figure 1.1. Two lines determine three or four regions.

Problem 1.1. What is the largest number of regions in a plane that three lines, four lines, or in general n lines determine? What is the smallest number of regions that n lines determine?

1.2 Diagonals and Triangulations

A triangle has no diagonals, a square has two and a convex pentagon has five, as in Figure 1.2. By "convex" we mean that it doesn't have any dents—see Figure 1.3. We define convex and diagonal after Figures 1.2 and 1.3. We often abbreviate "a polygon with n sides" as an *n-gon*. We'll call the corners of a polygon *vertices* and the sides *edges*, names that mathematicians use in broader contexts. (Your intuition of what a polygon is suffices for this book. However, the Appendix after Chapter 1 explores the thinking leading to a more formal definition.)

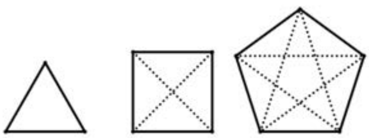

Figure 1.2. Diagonals of a triangle, a square, and a convex pentagon.

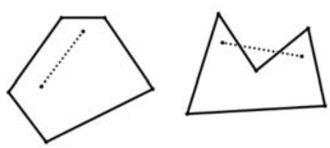

Figure 1.3. A convex shape and a nonconvex shape.

Definition. A set of points is *convex* provided for any two points P and Q in the set, the line segment $\overline{PQ}$ is entirely in the set. Otherwise, the set is *nonconvex*. We say a polygon is convex whenever its interior is a convex set.

Definition. For two vertices P and Q of a polygon, the line segment $\overline{PQ}$ is a *diagonal* provided it is not an edge of the polygon and, except for the endpoints P and Q, $\overline{PQ}$ is entirely in the interior of the polygon.

If a polygon is convex, whenever we pick two vertices not next to one another (not *adjacent*), the segment connecting them is a diagonal. So the number of diagonals depends only on the number of vertices, suggesting Problem 1.2. Once we ask about convex polygons, it is only natural to think about nonconvex ones, the topic of Problem 1.3.

Problem 1.2. How many diagonals does a convex polygon with n vertices have?

Problem 1.3. Do nonconvex polygons with at least four vertices always have some diagonals? If so, do all n-gons have a minimum number of diagonals (that may depend on n)?

From Figures 1.2 and 1.3, we can split the square into two triangles and each pentagon into three triangles, as in Figure 1.4.

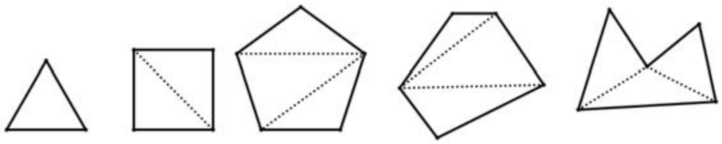

Figure 1.4. Triangulations of polygons from Figures 1.2 and 1.3.

Problem 1.4. Can we "triangulate" any polygon, as defined below, whether or not it is convex? Does every triangulation of an n-gon have the same number of triangles?

Definition. A *triangulation* of a polygon is a collection of triangles with nonoverlapping interiors, whose vertices are vertices of the polygon so that the triangles and their interiors together cover all of the polygon and its interior.

1.3 Distances and Points

Problem 1.5 (a). What is the smallest number of different (nonzero) distances n points on a line determine? What is the largest number? Can n points on a line determine any number of different distances between the smallest and largest numbers?

Problem 1.5 (b). Is there an arrangement of some number n of points in the plane that determine fewer distances than the smallest number of distances of n points in a line? What is the smallest number of distances 3, 4, 5, 6, 7, or 8 points in the plane can determine?

1.4 The Art Gallery Problem

In a museum or art gallery, guards (or nowadays cameras) are posted so that together they can see all the artwork to thwart thieves. If the museum or gallery were convex, one guard could survey all the art at once (provided there weren't people in the way). But real galleries and museums have far more complicated shapes. What is the fewest number of guards needed for any shaped polygons, based on the number of vertices? Figure 1.5 gives two examples. We state this more formally in Problem 1.6. The problem assumes the guards (the points G_i) can see in all directions, but not through the edges (walls) of the polygon.

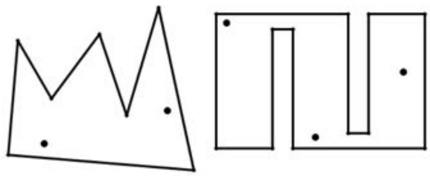

Figure 1.5. Art galleries needing two or three guards.

Definition. The points $G_1, G_2, \dots, G_k$ form a set of *guard points* for a region R provided that for all points P of R, there is some guard point G_i with the segment $\overline{G_iP}$ entirely contained in R.

Definition. The minimum number of guard points needed for any n-gon is g_n.

Problem 1.6. Explain why $g_4 = 1$, even though there are nonconvex quadrilaterals. Find g_5. Find a polygon with the fewest number of vertices needing two guard points. Repeat for three or more needed guard points. Look for a pattern for g_n.

1.5 Geometric Patterns

People throughout history and all over the world have delighted in designs, including colored patterns. Mathematicians have worked on classifying the possible types of patterns, colored or not. Figure 1.6 gives three two-colored patterns and a three-colored pattern for a regular hexagon that exhibit some sort of regularity in its coloring pattern, whereas the coloring of regions in Figure 1.7 seems irregular. (That is, it might be difficult to describe the pattern of how the colors are arranged in Figure 1.7.) In addition to looking for more such patterns in Problem 1.7 (a), Problem 1.7 (b) challenges you in a different way: exploring what we mean by regularity in coloring. The abstract nature of mathematics allows us to define concepts any way we want, but arbitrary definitions rarely prove worthwhile. Good definitions capture some key aspects of what we are studying and lead to interesting results.

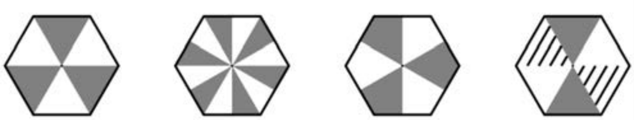

Figure 1.6. Regular colorings of a hexagon.

Figure 1.7. An irregular coloring of a hexagon.

Problem 1.7 (a). Design colored patterns for regular polygons with different numbers of vertices exhibiting regularity with two or more colors. What possible numbers of colors can such a coloring have in an n-gon?

Problem 1.7 (b). Describe what you mean by a pattern having a regular coloring.

1.6 Voronoi Diagrams

Given some fixed points in the plane, called *sites*, many people have investigated the regions of points closest to each of the fixed sites, starting at least with René Descartes (1596–1650). The resulting pictures, as in Figure 1.8, are called Voronoi diagrams, after the Ukrainian mathematician Georgy Feodoseevich Voronoi (1868–1908) who generalized them shortly before his death. Since then, scientists in disciplines ranging from astronomy to biology to engineering have found numerous applications of them.

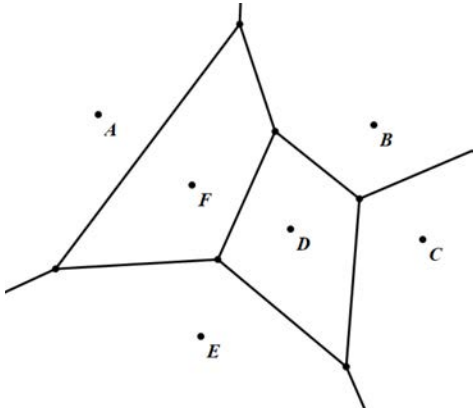

Figure 1.8. The Voronoi diagram for the six labelled points.

Definition. Given sites $S_1, S_2, \ldots, S_n$ in the plane the *Voronoi region* of S_i is the set of all points X so that for all sites S_k with $k \neq i$, $d(S_i, X) < d(S_k, X)$, where $d(X, Y)$ is the (usual Euclidean) distance between the points X and Y. The Voronoi regions together with their boundaries form the *Voronoi diagram* of the sites.

Problem 1.8. Find Voronoi diagrams when there are two, three, and four sites. What property or properties in terms of the given sites characterize the line segments, lines, and rays dividing the regions of a Voronoi diagram?

Remark. Along with puzzling about the preceding problems, you might be wondering "What about these geometric topics make them discrete?" "Discrete" means "separate," and in mathematics it is the opposite of continuous. For example, the lines in Problem 1 create separate regions, things we can count. Some of the other problems also asked you to count geometric items. In Problems 7 and 8 you didn't count anything, but you did color or construct discrete regions. Lines, planes, and many other geometric objects have continuous aspects, but these problems and almost all of those in Chapters 2, 3, and 4 focus on their discrete properties.

Appendix. What Is a Polygon?

Elementary school students develop the idea of a polygon from seeing lots of pictures of examples—and usually convex ones at that. Such an intuitive understanding of a polygon suffices for this book, but a formal definition is more involved. Mathematicians require careful definitions to prove results. They really don't like the phrase "the exception proves the rule" since they never want to allow a theorem to have any exceptions. So they look for potential *counterexamples*: things that are close to the intuitive idea of what they are trying to define, but fail in some way. The definition they craft needs to exclude the counterexamples while still including everything that fits their intuition. Let's use Figure 1.9 to compare the example of a polygon there with several designs that violate the common intuition of a polygon. After that we'll formulate a mathematical definition of a polygon.

A polygon is made up of vertices (corners) and edges (sides). Further, we expect a polygon to have a clear interior. The polygon on the left of Figure 1.9 has five vertices, five edges, and an interior.

The second drawing has four vertices and four edges, but two of the edges cross at a point that is not a vertex. Its interior appears as two triangular regions. A satisfactory definition of a polygon needs to eliminate examples with edges that intersect in a point other than a vertex. (In fact, we will use that property in the proof of Theorem 2.1 in the next chapter.) The third illustration adds a vertex at the intersection of the two previous edges, which seems to get around the

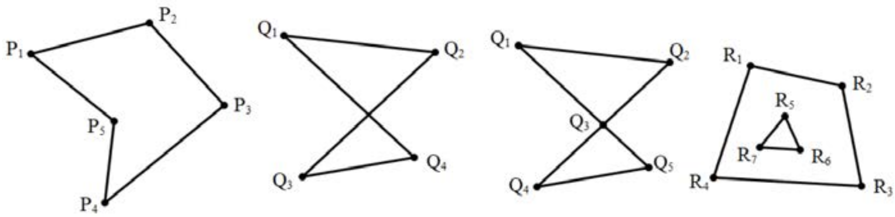

Figure 1.9. A polygon and three nonpolygons.

problem of crossing edges. But it is really two triangles sharing a vertex, rather than one polygon. So, we need to eliminate that option as well. In the design at the right, no edges cross, but it seems better described as a polygon with another polygon inside it, rather than just one polygon. (In Chapters 3 and 4, we'll consider polygons with holes like this.) Before we define a polygon, we define a segment of a line so that we can define an edge of a polygon.

Definition. The *segment* $\overline{PQ}$ is the set of points on the line through P and Q that are between P and Q as well as P and Q, which are called the *endpoints* of the segment.

Definition. An n-sided *polygon* (n-*gon*) is a set of n distinct points (called *vertices*) $P_1, P_2, \ldots, P_n$ in the plane and the n segments (called *edges*) $\overline{P_1P_2}$, $\overline{P_2P_3}$, $\ldots$, $\overline{P_{n-1}P_n}$, $\overline{P_nP_1}$ so that the intersection of two edges is either empty or is their common endpoint.

Check that the three counterexamples of Figure 1.9 fail to satisfy our definition of a polygon, but the pentagon of Figure 1.9 does satisfy it. Draw some figures you think should be polygons and some you think should not be polygons. Test which fit the definition above.

While it may seem obvious that polygons have an interior, mathematicians found a proof of this quite challenging. Camille Jordan (1838–1922) gave the first proof of the general theorem in 1887, which is now called the Jordan curve theorem in his honor. He showed that polygons and more generally "simple closed curves" separate the plane into an exterior and an interior. The theorem also shows how to tell whether two points are both inside or both outside the polygon (or curve). It does so by drawing the line segment between them and counting how many times it crosses the polygon. If it crosses an even number of times, they are both inside or both outside; if an odd number of times, one is inside and one is outside. It helps to have one of the points clearly outside or clearly inside. In Figure 1.10, the point C is clearly outside the polygon. Since the line segment $\overline{BC}$ crosses four edges of the polygon, point B is also outside. The segment $\overline{AC}$

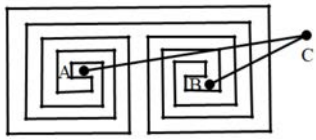

Figure 1.10. A polygon with some points inside and outside.

crosses nine edges, so A is inside the polygon. (See [7, 244–246] for more on this theorem.)

In addition to polygons, geometers have explored polygrams, also called star polygons. The best known is a pentagram, shown on the left of Figure 1.11. In polygrams, the edges cross in a set way, but otherwise they satisfy the definition of a polygon. (The numbering of the vertices of the pentagram enable the edges to fit the definition of a polygon, ignoring the edge crossing.) The second design of Figure 1.11 is not a polygram since it is made from two cycles of edges. No labeling of the vertices will make the edges fit the definition of a polygon, even ignoring the crossings. The second design is called a compound polygon. The last two drawings in Figure 1.11 are two different kinds of heptagrams. Note that the vertices of the two heptagrams are numbered differently so as to fit the labeling of the edges in the definition of a polygon. ("Hepta" is Greek for seven.)

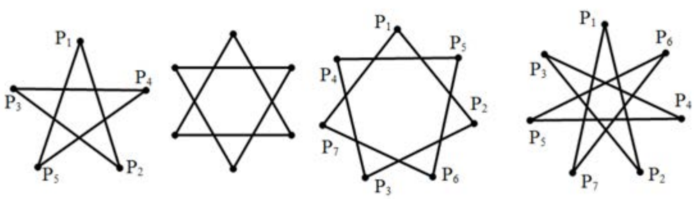

Figure 1.11. A pentagram, a compound polygon, and two heptagrams.

Exercise 1.1. Draw other polygrams and compound polygons whose vertices are the vertices of a regular n-gon. Look for conditions determining whether you get a polygram or a compound polygon. What conditions do we need to set on the crossings so that the second drawing in Figure 1.9 is not a polygram?

The answer to this exercise and the exercises in later chapters come after Chapter 4.

2

First Variations

Curiosity together with engaging problems propels mathematics forward. The following answers to the Chapter 1 problems suggest further problems to investigate. Each section follows the pattern of restating the related Chapter 1 problem or problems, answering them, and then questioning those ideas in at least one way. I encourage you also to try your hand at asking and answering your own variations of the problems.

2.1 Lines and Regions

Problem 1.1. (repeated) What is the largest number of regions in a plane that three lines, four lines, or in general n lines determine? What is the smallest number of regions that n lines determine?

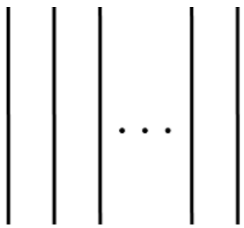

Figure 2.1. n parallel lines determine $n + 1$ regions

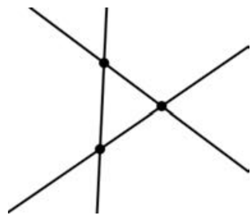

Figure 2.2. Three general lines determine seven regions.

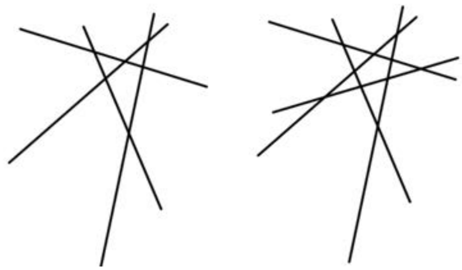

Figure 2.3. Regions determined by four and five lines in general position.

A set of n lines in the plane determines the fewest regions when they are all parallel to each other. They determine $n+1$ regions: $n-1$ slices between adjacent lines and the two half-planes on the ends, illustrated in Figure 2.1. In Figure 2.2, we see that three lines can determine seven regions. We say these lines are in *general position*—as opposed to the special cases of some parallel lines or three or more lines intersecting in the same point.

How can we find the maximum number of regions that n lines can determine? Draw some pictures to convince yourself first that to get the maximum number of regions, we need the lines to be in general position. Next, let's look at the values for $n = 1, 2, 3, 4$, and 5 lines to see whether we can find a pattern. Figure 2.3 gives examples of general cases for $n = 4$ and $n = 5$. Table 2.1 lists the maximum number of regions.

Table 2.1. Maximum number of regions for $n \leq 5$ lines.

lines	1	2	3	4	5
regions	2	4	7	11	16

The values in Table 2.1 suggest a pattern: the n^{th} line adds n more regions to what the maximum was before. This describes the maximum number of regions *recursively*, that is, values are determined using previous values. If we let $M(n)$ be the maximum number of regions determined by n lines, the pattern seems to be $M(n) = M(n-1)+n$. We will consider looking for a *closed form formula* for $M(n)$ shortly; that is, a formula depending only on n, the number of lines, not needing a previous value like $M(n-1)$. But first there is a more important mathematical question: why is the recursive pattern correct—why does the n^{th} line add at most n regions?

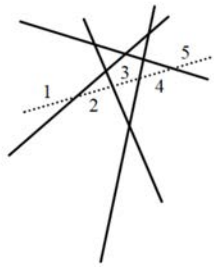

Figure 2.4. The pieces on the added line.

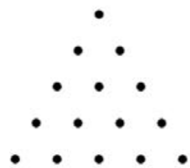

Figure 2.5. Triangular arrangement of points.

Suppose we have $n - 1$ lines making the maximum number of regions. In general position, the n^{th} line intersects each of the other $n - 1$ lines once. So, the new line is divided into n parts, including the two infinite ones, as labeled in Figure 2.4, where the dotted line is the new one. Note that each of the pieces of the new line divides one previous region in two, giving n new regions, as predicted. This argument also shows the importance of having the lines in general position. If the new line is parallel to any of the previous lines, we will miss an intersection and so fail to divide a region. Also, if the new line goes through the intersection of two or more previous lines, we will have fewer new regions. Thus $M(n)$ equals $M(n-1) + n$.

The numbers of regions in Table 2.1 are one more than the better known sequence of *triangular numbers* $1, 3, 6, 10, 15$. These numbers come from stacking points or balls in a triangular array. The n^{th} triangular number counts the number of points in the top n rows of Figure 2.5. The term "triangular numbers" and the even better known term "square numbers" have been around since the ancient Greeks, who also knew a formula for triangular numbers. (However, they had to describe the formula in words since our familiar algebraic notation is much more recent.) Figure 2.6 helps visualize the formula: the n^{th} triangular number is half of a "*rectangular number*," where one side of the rectangle is one longer than the other side. The formula for a rectangular number is $n(n + 1)$. In turn, the n^{th} triangular number is half of that, $\frac{n(n+1)}{2}$. Hence, the maximum number of regions in the plane determined by n lines is $M(n) = \frac{n(n+1)}{2} + 1 = \frac{n^2+n+2}{2}$.

One of the pleasures of mathematics comes after we answer a question—it is now "time to question the answer." That is, we modify the question to lead us in new directions and perhaps gain deeper insight. Here are several variations on Problem 1.1 for you to explore. You can invent your own variations of these problems and all of the others and try to answer them as well.

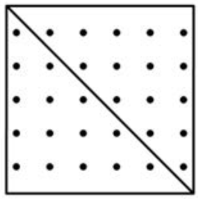

Figure 2.6. A rectangular number as the sum of two triangular numbers.

Problem 2.1 (a). By using some placement besides general position or all parallel lines can we have any number of regions between the minimum and maximum number of regions in the plane determined by n lines?

Problem 2.1 (b). What are the maximum and minimum number of regions in three-dimensional space determined by n planes?

Problem 2.1 (c). What are the maximum and minimum number of regions in the plane determined by n circles?

Problem 2.1 (d). What are the maximum number of regions in the plane determined by two convex polygons with j and k sides?

2.2 Diagonals and Triangulations

Problem 1.2. (repeated) How many diagonals does a convex polygon with n vertices have?

We can start as we did for Problem 1.1 about regions determined by lines—Table 2.2 lists the number of diagonals for different-sized convex polygons—and look for patterns. This problem illustrates one of the intriguing and beautiful aspects of mathematics: we can often solve a problem in different ways that yield the same answer.

Table 2.2. Number of diagonals in a convex polygon with n sides.

sides	3	4	5	6	7
diagonals	0	2	5	9	14

The difference in the number of diagonals goes up by one each time: $2-0=2$, $5-2=3$, $9-5=4$, etc. These differences showed up with Problem 1.1. Indeed, the values in Table 2.2 are one less than triangular numbers, although the n here doesn't match with the number k of numbers being added. Table 2.3 combines Table 2.2 with the corresponding triangular numbers.