

Chapter 4

Polygons

Although a line $\ell \subset \mathbb{R}^2$ can be treated as an infinite point set, it has a finite description. For instance, we may encode it by three coefficients $a, b, c \in \mathbb{R}$ with $(a, b) \neq (0, 0)$ and interpret it as all the points satisfying the equation $ax + by = c$. Actually all of the fundamental geometric objects that we mentioned in [Chapter 1](#) can be described by a constant number of parameters; hence they have constant *description complexity* or, informally, just *size*.¹

In this course we typically deal with objects that have unbounded size. Sometimes they are formed by merely aggregating constant-size objects. For instance, aggregation of points forms a finite point set. At other times we also demand additional structure beyond aggregation. Probably the most fundamental example is what we call a *polygon*. You probably learned this term in school, but what *is* a polygon precisely? Consider the examples shown in [Figure 4.1](#). Are these all polygons? If not, where would you draw the line?




Figure 4.1: *What is a polygon?*

4.1 Classes of Polygons

Ironically, there is no *the* right answer to the question, and there are different types of polygons. Usually, the sloppy term “polygon” refers to what we call a *simple polygon* defined below.

¹Unless specified otherwise, we always assume that the dimension is a small constant. If we work in high-dimensional space $\mathbb{R}^d$ where d varies, then their description complexity become $\Theta(d)$.

Definition 4.1. A **simple polygon** is a compact region $P \subset \mathbb{R}^2$ whose boundary is a simple closed curve $\partial P : [0, 1] \rightarrow \mathbb{R}^2$ consisting of finitely many consecutive line segments.

Out of the examples shown above only Polygon 4.1a is simple. For each of the remaining polygons the bounding segments do not make a simple closed curve.

When describing a simple polygon P it is sufficient to describe only its boundary ∂P .² As ∂P by definition consists of finitely many consecutive line segments, it can be encoded by a sequence $p_1, \dots, p_n$ of points, so that ∂P is formed by the line segments $\overline{p_1 p_2}, \overline{p_2 p_3}, \dots, \overline{p_{n-1} p_n}, \overline{p_n p_1}$. These points and segments are referred to as the *vertices* and the *edges* of the polygon, respectively. The set of vertices of a polygon P is denoted by $V(P)$, and the set of edges of P is denoted by $E(P)$.

Recall from Theorem 2.1 (Jordan curve theorem) that any simple closed curve separates the plane $\mathbb{R}^2$ into a (bounded) interior and an (unbounded) exterior. To prove this theorem in its full generality is surprisingly difficult. But for simple polygons the situation is easier, and in fact we can readily tell apart the interior from the exterior algorithmically, which we leave as an exercise.

Exercise 4.2. Describe an algorithm to decide whether a point lies inside or outside of a simple polygon. More precisely, given a simple polygon $P \subset \mathbb{R}^2$ as a list of its vertices $(v_1, v_2, \dots, v_n)$ in counterclockwise order and a query point $q \in \mathbb{R}^2$, decide whether q is inside P , on the boundary of P , or outside. The runtime of your algorithm should be $O(n)$.

There are good reasons to ask for the boundary of a polygon to form a simple curve: For instance, in the example depicted in Figure 4.1b there are several regions for which it is completely unclear whether they should belong to the interior or to the exterior of the polygon. A similar problem arises for the interior regions in Figure 4.1f. But there are more general classes of polygons that some of the remaining examples fall into. We will discuss only one such class here. It comprises polygons like the one from Figure 4.1d.

Definition 4.3. A region $P \subset \mathbb{R}^2$ is a **simple polygon with holes** if it can be described as $P = F \setminus \bigcup_{H \in \mathcal{H}} H^\circ$, where $\mathcal{H}$ is a finite collection of pairwise disjoint simple polygons (called *holes*) and F is a simple polygon for which $F^\circ \supset \bigcup_{H \in \mathcal{H}} H$.

The way we define them through the notion of simple polygons makes a trichotomy immediate, just as for simple polygons: Every point in the plane can be either inside, on the boundary, or outside of P .

4.2 Polygon Triangulation

Topologically speaking, a simple polygon is nothing but a disk and thus a very elementary object. But geometrically a simple polygon can be—as if mocking the label we attached

²In fact this is not obvious. To see the subtlety, consider a polygon P and the point set $Q := P^\circ \cup (Q^2 \cap \partial P)$. Note that $\partial P = \partial Q$, so it is possible to recover different sets out from the same boundary. The catch here is that Q is not compact. One can show that we can recover only one *compact* set out from a given boundary, and that is why we need compactness in the definition of a simple polygon.




Figure 4.2: A simple (?) polygon.

to it—pretty complicated in shape, see Figure 4.2 for an example. While it has a succinct one-dimensional representation as the sequence of the boundary vertices, we often want to work with a more structured representation that retains the two-dimensional shape. For instance, computing the area of a general simple polygon is not so straightforward out of a one-dimensional representation. But if we manage to represent the polygon as a disjoint union of simpler geometric shapes such as triangles, rectangles or trapezoids, then its area simply sums up all the area of the individual shapes, which are easy to compute. This motivates the definition of a triangulation.

Definition 4.4. A triangulation of a simple polygon P with vertex set $V(P)$ is a collection $\mathcal{T}$ of triangles, such that

- (1) $P = \bigcup_{T \in \mathcal{T}} T$;
- (2) $V(P) = \bigcup_{T \in \mathcal{T}} V(T)$; and
- (3) for every distinct pair $T_1, T_2 \in \mathcal{T}$, the intersection $T_1 \cap T_2$ is either a common vertex, a common edge, or empty.

Exercise 4.5. Show that each condition in Definition 4.4 is necessary in the following sense: Give an example of a non-triangulation that would form a triangulation if the condition was omitted. Is the definition equivalent if (3) is replaced by $T_1^\circ \cap T_2^\circ = \emptyset$, for every distinct pair $T_1, T_2 \in \mathcal{T}$?

Beyond area computation, triangulations are incredibly useful in planar geometry. The significance roots in the fact that every simple polygon admits a triangulation.

Theorem 4.6. Every simple polygon has a triangulation.

Proof. Let P be a simple polygon on n vertices. We prove the statement by induction on n . For $n = 3$, the polygon P is a triangle, which forms a triangulation by itself. For $n > 3$, consider the lexicographically smallest vertex v of P . That is, among all vertices with the smallest x -coordinate we pick the one with the smallest y -coordinate. Let vertices u and w be the predecessor and successor of v along a counterclockwise traversal of ∂P , respectively, so that P is locally to the left of $u \rightarrow v \rightarrow w$. By choice of v , the walk $u \rightarrow v \rightarrow w$ forms a strict left turn. Hence $T := uvw$ forms a triangle that lies in P . We distinguish two cases.

Case 1: $\text{relint}(\overline{uw}) \subset P^\circ$ (Figure 4.3a). That is, the segment $\overline{uw}$ except for its two endpoints is fully contained in P° . Hence P can be split into two interior-disjoint simple polygons: the triangle T and a polygon $P' := (P \setminus T) \cup \overline{uw}$ on $n - 1$ vertices. By the inductive hypothesis, the polygon P' admits a triangulation. This triangulation together with T yields a triangulation of P .




Figure 4.3: Cases in the proof of Theorem 4.6.

Case 2: $\text{relint}(\overline{uw}) \not\subset P^\circ$ (Figure 4.3b). Then some point from ∂P must be in T° or on $\overline{uw}$. As ∂P is composed of line segments, some vertex of P must be in T° or on $\overline{uw}$. Among all such vertices select p to maximize the distance to the line uw . Imagine a line through p that is parallel to uw . This line together with uv, vw bounds a triangle that does not contain any part of ∂P in its interior. It follows that $\text{relint}(\overline{vp}) \subset P^\circ$. Therefore, the segment vp splits P into two interior-disjoint polygons P_1 and P_2 on fewer than n vertices each (the vertex u does not appear in one of them, whereas w does not appear in the other). By the inductive hypothesis, both P_1 and P_2 have triangulations, and their union yields a triangulation of P . $\square$

Exercise 4.7. In the proof of Theorem 4.6, would the argument in Case 2 also work if the point p was chosen to be a vertex of P in T° that minimizes the Euclidean distance to v ?

The configuration from Case 1 above is called an *ear*: three consecutive vertices u, v, w of a simple polygon P such that the relative interior of $\overline{uw}$ lies in P° . In fact, we could have skipped the analysis for Case 2 by referring to the following theorem.

Theorem 4.8 (Meisters [13, 14]). *Every simple polygon with $n \geq 4$ vertices has two ears A and B such that $A^\circ \cap B^\circ = \emptyset$.*

But conversely, knowing Theorem 4.6 and the theorem below, we can recover Theorem 4.8 as a direct consequence.

Theorem 4.9. *Every triangulation of a simple polygon P with $n \geq 4$ vertices contains at least two triangles that serve as ears in P .*

Exercise 4.10. Prove [Theorem 4.9](#).

Exercise 4.11. Let P be a simple polygon with vertices $v_1, v_2, \dots, v_n$ in counterclockwise order, where v_i has coordinates (x_i, y_i) . Show that the area of P is

$$\frac{1}{2} \sum_{i=1}^n x_i y_{i+1} - x_{i+1} y_i,$$

where we agree that $(x_{n+1}, y_{n+1}) = (x_1, y_1)$.

A triangulation naturally induces a plane straight-line graph, whose vertices are the polygon vertices and whose edges come from the sides of triangles. The number of edges and triangles (faces) are completely determined by the number of vertices, as the following simple lemma shows.

Lemma 4.12. Every triangulation of a simple polygon on $n \geq 3$ vertices consists of $n - 2$ triangles and $2n - 3$ edges.

Proof. We prove the statement by induction on n . For $n = 3$, a triangulation has $n - 2 = 1$ triangle and $2n - 3 = 3$ edges. Thus, the statement holds.

In the general case, take any triangulation T of a simple polygon P on $n \geq 4$ vertices. Let X be a triangle of T that serves as an ear of P , which exists by [Theorem 4.9](#). Then $T - X$ is a triangulation of a simple polygon on $n - 1$ vertices. By the inductive hypothesis, the triangulation $T - X$ consists of $(n - 1) - 2 = n - 3$ triangles and $2(n - 1) - 3 = 2n - 5$ edges. Putting X back adds one more triangle and two more edges. Therefore, the triangulation T consists of $n - 3 + 1 = n - 2$ triangles and $2n - 5 + 2 = 2n - 3$ edges. $\square$

Tetrahedralizations in $\mathbb{R}^3$. The universal presence of triangulations is something particular about the plane: The natural generalization of [Theorem 4.6](#) to dimension three and higher does not hold. Here we discuss the issue in $\mathbb{R}^3$.

A simple polygon is a topological disk in $\mathbb{R}^2$ that is locally bounded by patches of lines. The corresponding term in $\mathbb{R}^3$ is a *polyhedron*, which can be informally defined via a literal translation of the previous sentence: it is a topologically ball that is locally bounded by patches of planes. A triangle in $\mathbb{R}^2$ corresponds to a tetrahedron in $\mathbb{R}^3$ and a *tetrahedralization* is a nice partition into tetrahedra. Being “nice” means that the union of the tetrahedra covers the object, the vertices of the tetrahedra are vertices of the polyhedron, and any two distinct tetrahedra intersect in either a common triangular face, a common edge, or a common vertex, or not at all.³

Unfortunately, there are polyhedra in $\mathbb{R}^3$ that do not admit a tetrahedralization. The following construction is due to Schönhardt [18]. It is based on a triangular prism, that is, two congruent triangles placed in parallel planes, where the corresponding sides of both triangles are connected by a rectangle ([Figure 4.4a](#)). We slightly rotate one

³These “nice” subdivisions can be defined in an abstract combinatorial setting, where they are called *simplicial complexes*.

triangle within its plane, thus twisting the prism. As a consequence, the rectangles are dented inward along their diagonal in direction of the rotation, and are no longer plane. (Figure 4.4b). The other diagonals of the (former) rectangles—labeled ab' , bc' , and




Figure 4.4: *The Schönhardt polyhedron cannot be subdivided into tetrahedra without adding new vertices.*

ca' —are now epigonals, that is, they lie in the exterior of the polyhedron. Since these epigonals are the only missing edges between the vertices, there is no way to add edges to form a tetrahedron for a subdivision. Clearly the polyhedron is not a tetrahedron by itself, and so we conclude that it does not admit a tetrahedralization. Actually, it is NP-complete to decide whether a non-convex polyhedron has a tetrahedralization [15]. However, if adding new vertices—which are called Steiner vertices—is allowed, then a tetrahedralization is possible, both in this example and in general.

Even if a tetrahedralization of a polyhedron exists, there is another significant difference to polygons in $\mathbb{R}^2$. While the number of triangles in a triangulation of a polygon depends only on the number of vertices, the number of tetrahedra in two different tetrahedralizations of the same polyhedron may be different. See Figure 4.5 for a simple example of a polyhedron that has tetrahedralizations with two or three tetrahedra. Deciding whether a convex polyhedron has a tetrahedralization with at most a given number of tetrahedra is NP-complete [6].

Exercise 4.13. *Characterize all possible tetrahedralizations of the three-dimensional cube.*

Triangulation algorithms. Knowing that a triangulation exists is nice, but it is even nicer to know that it can also be constructed efficiently.

Exercise 4.14. *Convert Theorem 4.6 into an $O(n^2)$ time algorithm to construct a triangulation for a given simple polygon with n vertices.*

The runtime of this straightforward application of Theorem 4.6 is not optimal. For those who are interested in more efficient algorithms, please refer to Appendices A and C.




Figure 4.5: *Two tetrahedralizations of the same polyhedron, a triangular bipyramid. The left partition uses two polyhedra; both the top vertex t and the bottom vertex b belong to only one tetrahedron. The right partition uses three polyhedra that all share the dashed diagonal bt .*

The best (in terms of worst-case runtime) algorithm known due to Chazelle [7] computes a triangulation in linear time, but it is very complicated. There is also a somewhat simpler randomized algorithm in expected linear time [4]. We will not cover either of them in the notes. It remains open whether there exists a simple (which is not really well-defined, except that Chazelle’s Algorithm does not qualify) deterministic linear time algorithm to triangulate a simple polygon [10].

Polygons with holes. It is interesting to note that the complexity of the triangulation problem changes to $\Theta(n \log n)$, if the polygon may contain holes [5]. This means that (1) there is an algorithm to construct a triangulation for a given simple polygon with holes on n vertices (counting both the vertices on the outer boundary and on the holes’ boundaries) in $O(n \log n)$ time; and (2) there is a lower bound of $\Omega(n \log n)$ operations in any model of computation that is subject to the same lower bound for comparison-based sorting. This difference in complexity is a very common pattern: There are many problems that are (sometimes much) harder for simple polygons with holes than for simple polygons. So maybe the term “simple” has some justification, after all...

General triangle covers. What if we drop the “niceness” conditions for triangulations and just want to describe a given simple polygon as a union of triangles? It turns out this is a rather drastic change. For instance, it is unlikely that we can efficiently find an optimal/minimal description of this type: Christ has shown [8] that it is NP-hard to decide whether for a simple polygon P with n vertices and a positive integer k , there exists a set of at most k triangles whose union is P . In fact, the problem is not even known to be in NP, because it is not clear whether the coordinates of solutions can always be encoded succinctly.

4.3 The Art Gallery Problem

In 1973 Victor Klee posed the following question: “How many guards are necessary, and how many are sufficient to patrol the paintings and works of art in an art gallery with n walls?” Thinking in geometric terms, we may model “an art gallery with n walls” as a simple polygon P bounded by n edges, hence also n vertices. And a guard is modeled as a point $g \in P$ where he stands. The edges are opaque and prevent one from seeing what is behind, thus g watches over all the points p for which the line segment $\overline{gp}$ lies completely in P ; see Figure 4.6. The task is then to place as least points as possible so that the entire P is being watched.




Figure 4.6: The region that a guard g can observe.

It is not hard to see that $\lfloor n/3 \rfloor$ guards are necessary in some cases.

Exercise 4.15. Describe a family $(P_n)_{n \geq 3}$ of simple polygons such that P_n has n vertices and requires at least $\lfloor n/3 \rfloor$ guards.

What is more surprising: $\lfloor n/3 \rfloor$ guards are always sufficient for all P with n vertices. Chvátal [9] was the first to show it, but then Fisk [11] gave a much simpler proof using—as you may have guessed—triangulations. Fisk’s proof was considered so beautiful that was selected in “Proofs from THE BOOK” [3]. It is based on the following lemma.

Lemma 4.16. Every triangulation of a simple polygon is 3-colorable. That is, each vertex can be assigned one of three colors so that adjacent vertices receive different colors.

Proof. Induction on n . For $n = 3$ the statement is obvious. For $n > 3$, by Theorem 4.9 the triangulation contains an ear uvw . Cutting off the ear creates a triangulation of a polygon on $n - 1$ vertices, which by the inductive hypothesis admits a proper 3-coloring. Now whichever two colors the vertices u and w receive in this coloring, there remains a spare color for v . $\square$

Theorem 4.17 (Fisk [11]). Every simple polygon with n vertices can be guarded using at most $\lfloor n/3 \rfloor$ guards.




Figure 4.7: *A triangulation of a simple polygon on 17 vertices and a 3-coloring of it. The orange vertices form the smallest color class and guard the polygon using $5 \leq \lfloor 17/3 \rfloor$ guards.*

Proof. Fix a triangulation of the polygon and then a 3-coloring of the vertices, as ensured by [Lemma 4.16](#). Place a guard at each vertex of the smallest color class, which clearly amounts to at most $\lfloor n/3 \rfloor$ vertices. As any point p in the polygon lands in some triangle T and exactly one of T 's vertices has the selected color, the point p is watched by that vertex. Hence the whole polygon is guarded. $\square$

4.4 Optimal Guarding

While [Exercise 4.15](#) shows that the bound in [Theorem 4.17](#) is tight in general, it is easy to see that Fisk's method does not necessarily minimize the number of guards. Also, we do not have to restrict ourselves to place the guards at vertices only, but can rather place them anywhere on the boundary or even in the interior of the polygon. In all these cases, we can ask for the minimum number of guards required to guard a given polygon P . These problems have been shown to be NP-hard by Lee and Lin [\[12\]](#) already in the 1980s. Actually, if the guards are not constrained to lie on vertices, it is not even clear whether the corresponding decision problem is in NP. In fact, recent results by Abrahamsen et al. suggest the opposite. In the remainder of this section we will briefly discuss some of these results.

Recall that a decision problem is in NP if for any “yes” instance, one can present a *certificate* that allows polynomial-time verification of the “yes” status.⁴ In our context, if we restrict the guards to be on vertices, a natural certificate is the set of vertices where the guards stand. It allows us to verify that the guards indeed watch the entire polygon and that the number of guards is within the specified limit.

However, if we drop the restriction, a natural certificate would be the coordinates of the guards. Since no more than $\lfloor n/3 \rfloor$ guards are required, this seems a reasonable certificate. But what if the number of bits needed to explicitly represent these coordinates are exponential in input size? One might be tempted to think that any reasonable guard can be placed at an intersection point of some lines that are defined by polygon vertices.

⁴And of course, for any “no” instance there should not be any fake certificate that tricks the verification...




Figure 4.8: *To guard this polygon with three guards, there must be one guard on each of the green dashed segments. The middle guard g_m must be to the left of the blue curve, to the right of the red curve, and on the dashed green line. The intersection point of these three curves has irrational coordinates.*

Alas, in general this is not correct: some guards with irrational coordinates may be required, even if all vertices of P have integral coordinates! This surprising result has been presented in 2017 and we will sketch its main ideas, referring to the paper by Abrahamsen, Adamaszek, and Miltzow [1] for more details and the exact construction.

Consider the polygon shown in Figure 4.8, which consists of a main rectangular region with triangular, rectangular, and trapezoidal regions attached. On the one hand, it can be watched by three guards. On the other hand, we will argue that, if this polygon is guarded with less than four guards, at least one of the guards has an irrational coordinate. The polygon contains three pairs of triangular regions with the following structure. Each pair is connected by a green dashed segment in the figure. This segment contains one edge of each of the two triangles and separates their interiors. Hence, a single guard that sees both of these triangles has to be placed on this separating segment. Further, there is no other point that can guard two of these six triangles. Therefore, if we have only three guards, each of them must be placed on one of these three disjoint segments. The small rectangular regions to the left, top, and bottom outside the main rectangular region further constrain the positions of the guards along these segments.

Let the guards be g_l , g_m , and g_r , as in the figure. The guard g_l cannot see all the points inside the left two trapezoidal regions, and thus g_m has to be placed appropriately. For each position of g_l on its segment, we get a unique rightmost position on which a second guard can be placed to guard the two trapezoids. The union of these points defines an arc that is a segment of a quadratic curve (the roots of a quadratic polynomial). We get an analogous curve for g_r and the two trapezoids attached to the right. By a careful choice of the vertex coordinates, these two curves cross at a point p that also lies on the segment for the guard g_m and has irrational coordinates. It then follows from a detailed argument (see [1]) that p is the only feasible placement of g_m . Let us point out that the choice of the vertex coordinates to achieve this is far from trivial. For example, there can only be a single line defined by two points with rational coordinates that passes through p , and this is the line on which the guard g_m is constrained to lie on.

Exercise 4.18. Let P be a polygon with vertices on the integer grid, and let g be a point inside that polygon with at least one irrational coordinate. Show that there can be at most one diagonal of P passing through g .

Nevertheless, the sketched construction leads to the following result.

Theorem 4.19 (Abrahamsen et al. [1]). For any k , there is a simple polygon P with integer vertex coordinates such that P can be guarded by $3k$ guards, while a guard set having only rational coordinates requires $4k$ guards.

Abrahamsen, Adamaszek, and Miltzow [2] showed recently that the art gallery problem is actually complete in a complexity class called $\exists\mathbb{R}$. The *existential theory of the reals* (see [16, 17] for details) is the set of true sentences of the form $\exists x_1, \dots, x_n \in \mathbb{R} : \phi(x_1, \dots, x_n)$ for a quantifier-free Boolean formula ϕ without negation⁵ that can use the constants 0 and 1 as well as the operators $+$, $*$, and $<$. For example, $\exists x, y : (x < y) \wedge (x * y < 1 + 1)$ is such a formula. A problem is in the complexity class $\exists\mathbb{R}$ if it allows for a polynomial-time reduction to the problem of deciding such formulas, and it is complete if in addition every problem in $\exists\mathbb{R}$ can be reduced to it in polynomial time.

Exercise 4.20. Show that $\text{NP} \subseteq \exists\mathbb{R}$.

For the art gallery problem, the result by Abrahamsen et al. [2] implies that the coordinates of an optimal guard set may be doubly-exponential in the input size. This statement does not exclude the possibility of a more concise, implicit way to express the existence of an optimal solution. However, if we found such a way, then this would imply that the art gallery problem is in NP , which, in turn, would imply $\text{NP} = \exists\mathbb{R}$.

Questions

11. What is a simple polygon/a simple polygon with holes? Explain the definitions and provide some examples of members and non-members of the respective classes. For a given polygon you should be able to tell which of these classes it belongs to or does not belong to and provide justifications.
12. What is a closed/open/bounded set in $\mathbb{R}^d$? What is the interior/closure of a point set? Explain the definitions and provide some illustrative examples. For a given set you should be able to argue which of the aforementioned properties it possesses.
13. What is a triangulation of a simple polygon? Does it always exist? Explain the definition and provide some illustrative examples. Present the proof of Theorem 4.6 in detail.

⁵If we also allowed negation (and hence also the relations $\leq$ and $=$), then it would be possible to express statements like $\exists x : x * x = 1 + 1$ that have only irrational solutions, which is impossible to achieve by using only strict inequalities $<$. Interestingly, however, the complexity class $\exists\mathbb{R}$, when defined in terms of these more expressive formulas, would remain the same, see [17].

14. *How many points are needed to guard a simple polygon?* Present the proofs of [Theorem 4.9](#), [Lemma 4.16](#), and [Theorem 4.17](#) in detail.
15. *How about higher dimensional generalizations? Can every polyhedron in $\mathbb{R}^3$ be nicely subdivided into tetrahedra?* Explain Schönhardt's construction.
16. (This topic was not covered in HS25 and therefore the question will not be asked in the exam.) *Is there a succinct representation for optimal guard placements?* State [Theorem 4.19](#) and sketch the construction.

References

- [1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow, [Irrational guards are sometimes needed](#). In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pp. 3:1–3:15, 2017.
- [2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow, [The art gallery problem is \$\exists\mathbb{R}\$ -complete](#). In *Proc. 50th Annu. ACM Sympos. Theory Comput.*, pp. 65–73, 2018.
- [3] Martin Aigner and Günter M. Ziegler, [Proofs from THE BOOK](#), Springer, Berlin, 3rd edn., 2003.
- [4] Nancy M. Amato, Michael T. Goodrich, and Edgar A. Ramos, [A randomized algorithm for triangulating a simple polygon in linear time](#). *Discrete Comput. Geom.*, 26/2, (2001), 245–265.
- [5] Takao Asano, Tetsuo Asano, and Ron Y. Pinter, [Polygon triangulation: efficiency and minimality](#). *J. Algorithms*, 7/2, (1986), 221–231.
- [6] Alexander Below, Jesús A. De Loera, and Jürgen Richter-Gebert, [The complexity of finding small triangulations of convex 3-polytopes](#). *J. Algorithms*, 50/2, (2004), 134–167.
- [7] Bernard Chazelle, [Triangulating a simple polygon in linear time](#). *Discrete Comput. Geom.*, 6/5, (1991), 485–524.
- [8] Tobias Christ, [Beyond triangulation: covering polygons with triangles](#). In *Proc. 12th Algorithms and Data Struct. Sympos.*, vol. 6844 of *Lecture Notes Comput. Sci.*, pp. 231–242, Springer, 2011.
- [9] Václav Chvátal, [A combinatorial theorem in plane geometry](#). *J. Combin. Theory Ser. B*, 18/1, (1975), 39–41.
- [10] Erik D. Demaine, Joseph S. B. Mitchell, and Joseph O'Rourke, The Open Problems Project, Problem #10. <https://topp.openproblem.net/p10>.

- [11] Steve Fisk, [A short proof of Chvátal's watchman theorem](#). *J. Combin. Theory Ser. B*, 24/3, (1978), 374.
- [12] Der-Tsai Lee and Arthur K. Lin, [Computational complexity of art gallery problems](#). *IEEE Trans. Inform. Theory*, 32/2, (1986), 276–282.
- [13] Gary H. Meisters, [Polygons have ears](#). *Amer. Math. Monthly*, 82/6, (1975), 648–651.
- [14] Gary H. Meisters, [Principal vertices, exposed points, and ears](#). *Amer. Math. Monthly*, 87/4, (1980), 284–285.
- [15] Jim Ruppert and Raimund Seidel, [On the difficulty of triangulating three-dimensional nonconvex polyhedra](#). *Discrete Comput. Geom.*, 7, (1992), 227–253.
- [16] Marcus Schaefer, [Complexity of some geometric and topological problems](#). In *Proc. 17th Int. Sympos. Graph Drawing (GD 2009)*, pp. 334–344, 2009.
- [17] Marcus Schaefer and Daniel Štefankovič, [Fixed Points, Nash Equilibria, and the Existential Theory of the Reals](#). *Theory Comput. Syst.*, 60, (2017), 172–193.
- [18] Erich Schönhardt, [Über die Zerlegung von Dreieckspolyedern in Tetraeder](#). *Math. Ann.*, 98, (1928), 309–312.