Chapter 1

Art Gallery Problems

1.1 Introduction

Illumination problems have always been a popular topic of study in mathematics. For example, it is well known that the boundary of any smooth compact convex S set on the plane can always be illuminated using three light sources located in the complement of S. An easy proof of this can be obtained by enclosing a smooth convex set within a triangle, then placing a light at each vertex of T; see Figure 1.1.

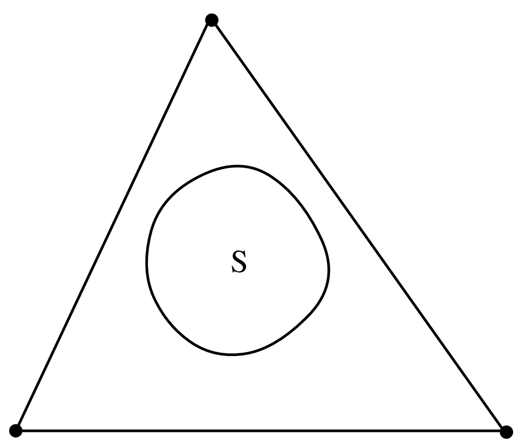

Figure 1.1: Three lights suffice to illuminate a compact convex set.

One famous—and until recently open—problem on illumination is attributed to Ernst Strauss (see E. G. Strauss and V. Klee [83]), who, in the early fifties posed the following problem:

Suppose that we live in a two-dimensional room whose walls form a simple closed polygon P and each wall is a mirror.

1. Is it true that if we place a light at any point of P, all of P will be illuminated using reflected rays as well as direct rays?

2. Is there necessarily a point from which a single light source will illuminate the entire room using reflected rays as well as direct rays?

The first part of Strauss's problem was recently proved to be false by G. W. Tokarsky [127]. Tokarsky's proof is surprisingly simple, using basic concepts of geometry that are easily understandable. We refer the interested reader to Tokarsky's original manuscript which is clearly written and a pleasure to read. The second part of Strauss's conjecture, though, remains open. It would be nice if a "simple" proof for it could be obtained.

More closely related to our topic of interest here is a question posed by V. Klee during a conference in Stanford in August 1976. Klee's question was:

> How many guards are always sufficient to guard any polygon with n vertices?

Soon after, V. Chvátal established what has become known as *Chvátal's Art Gallery Theorem*, namely: that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and occasionally necessary to cover a simple polygon with n vertices [28].

Since the publication of this original result, a tremendous amount of research on illumination or Art Gallery problems has been carried out by mathematicians and computer scientists. In 1987, J. O'Rourke [104] published *Art Gallery Theorems and Algorithms*, the first book dedicated solely to the study of illumination problems of polygons on the plane. The publication of O'Rourke's book further fueled the study of Art Gallery type problems, and many variations to the original Art Gallery Theorem have since been studied. Two thorough survey papers have also been written on this subject, one in 1992 by T. Shermer [115], and a second one in 1996 by J. Urrutia [131]. In this book, we present most of the material known to us on Art Gallery or illumination problems. Following J. O'Rourke's lead, we have tried to write this book in a manner that makes the material presented here accessible to a wide audience. The only background required for this book is basic knowledge of algorithms, graph theory and data structures. This should make this book both accessible to senior undergraduate students and suitable for courses at the master's level. Numerous open problems are presented, making the book also useful to researchers working in this area.

1.2 Basic terminology

A *polygon* P is an ordered sequence of points $p_1, \ldots, p_n$, $n \geq 3$, called the vertices of P together with the set of line segments joining p_i to p_{i+1}, $i = 1, \ldots, n-1$ and p_n to p_1, called the edges of P. P is called *simple* if any two non-consecutive edges do not intersect. A simple polygon divides the plane into two regions, an unbounded one called the exterior region and a bounded one, the interior. Henceforth to simplify our presentation, the term polygon will be used to denote simple polygons together with their interior.

A polygon is called *orthogonal* if all its edges are parallel to either the x-axis or the y-axis.

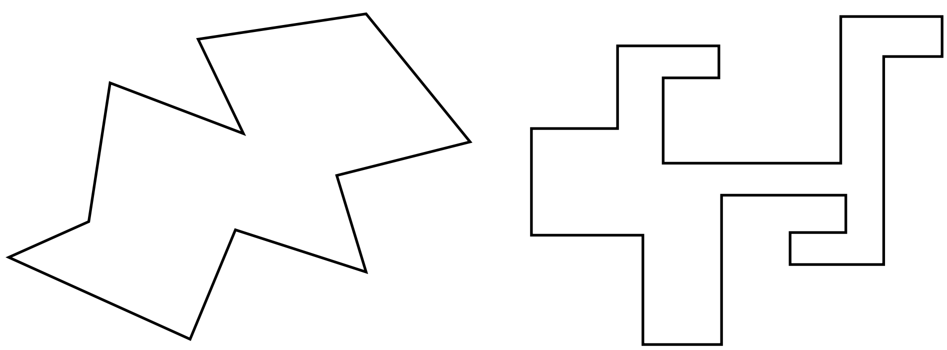

Figure 1.2: A simple and an orthogonal polygon.

A *graph* G consists of a set of elements V called the vertices of G, together with a set E of pairs of vertices of G called the edges of G. Two vertices u and v of G are called *adjacent* if the pair $\{u, v\}$ is an element of E. The *degree* of a vertex v of G is the number of vertices of G adjacent to v. A graph G is planar if it can be drawn on the plane in such a way that its vertices are represented by points on the plane, and each edge $\{u, v\}$ of G is represented by a simple curve joining points representing u and v. Moreover, two edges of G may only intersect at their endpoints. In Figure 1.3 we show two graphs; the first one is planar, and the second one is not.

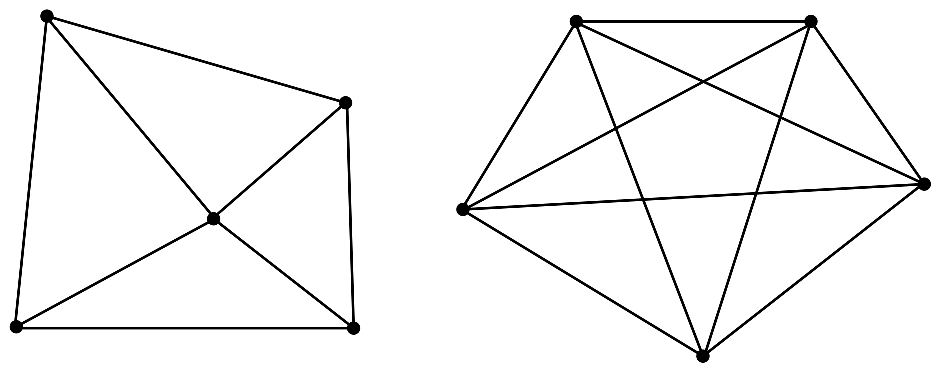

Figure 1.3: Two graphs with five vertices.

A *path* of a graph G is a sequence of distinct vertices $v_1, \ldots, v_k$ such that v_i and v_{i+1} are adjacent in G, $i = 1, \ldots, k-1$, $k \geq 2$. A *cycle* of G is a path $v_1, \ldots, v_k$ together with the edge $\{v_k, v_1\}$, $k \geq 3$. A graph $G(V, E)$ is called connected if for every pair of vertices u and v of $G(V, E)$, there exists a path $u = v_1, \ldots, v_k = v$ starting at u and ending at v; otherwise $G(V, E)$ is called disconnected. A graph G is called a *tree* if it is connected and contains no cycles.

A *triangulation* T of a polygon P is a partitioning of P into a set of triangles with pairwise disjoint interiors in such a way that the edges of those triangles are either edges or diagonals of P joining pairs of vertices.

The following result is easy to prove:

Theorem 1.2.1 *Any polygon can be triangulated; moreover any triangulation of a polygon P with n vertices contains exactly $n-2$ triangles.*

Triangulations of polygons play a central role in the study of art gallery problems. The problem of finding efficient algorithms to triangulate polygons has received much attention in computational geometry. In 1978, Garey, Johnson, Preparata and Tarjan [60] obtained the first $O(n \ln n)$ time triangulation algorithm. This result was improved in 1988 by Tarjan and van Wyk to $O(n \ln \ln n)$ [126]. Finally Chazelle [20] proved in 1990:

Theorem 1.2.2 *There is a linear time algorithm to triangulate simple polygons.*

The algorithmic details of Chazelle's triangulation algorithm are beyond the scope of this book, and will not be studied here.

Given a triangulation T of a polygon P, we can define a graph $GT(P)$ such that the vertices of $GT(P)$ are the vertices of P, and two vertices of $GT(P)$ are adjacent if they are connected by an edge of T.

The dual of a triangulation T of a polygon P is obtained by placing a vertex in the interior of each triangle of T, and connecting two vertices if their corresponding triangles share a common edge in T. It is easy to see that the dual graph of a triangulation of a polygon P is always a tree. In Figure 1.4 we show a polygon together with a triangulation T and its dual.

A *coloring* of the vertices of a graph G is an assignment of colors to its vertices in such a way that adjacent vertices receive different colors. The chromatic number of a graph G is the smallest integer k such that a coloring of G exists. In this case, we say that G has chromatic number k. The following result can be easily proved:

Theorem 1.2.3 *Let P be a polygon and T a triangulation of P. Then $GT(P)$ has chromatic number 3.*

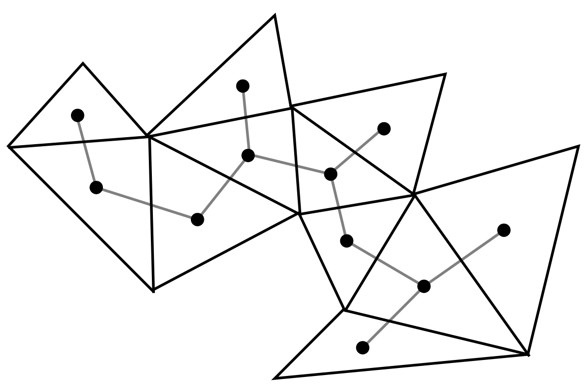

Figure 1.4: A triangulation of a polygon together with its dual.

Given two points p and q of a polygon P, we say that p is *visible* from q if the line segment joining p to q is totally contained in P.

A collection H of points of P *illuminates* or *guards* P if every point u of P is visible from a point p in H. The term *illuminates* follows the notion that if we place a light source that emits light *in all directions* on each element of H, P is totally illuminated. The use of the term *guard* follows the notion that if we station a guard at each element of H, all of P is guarded. The reader may easily verify that to illuminate the orthogonal polygon in Figure 1.2 we need four lights. We remark at this point that the terms *illumination* and *guarding* will be used interchangeably in this manuscript. Our choice of whether to "guard" or "illuminate" an object depends mainly on the term used in the original paper in which a particular result was proved.

The visibility graph $VG(P)$ of P is the graph whose vertex set is the set of vertices of P, two vertices u and v being adjacent in $VG(P)$ if they are visible in P. Visibility graphs were first introduced by Avis and ElGindy [6].

Avis and ElGindy gave an optimal (worst case) $O(n^2)$ time algorithm to compute the visibility graph of a polygon. Their result was later improved by Hershberger [72], who gave an $O(\mid E \mid)$ time algorithm, where $\mid E \mid$ is the number of edges of the visibility graph of P. One of the most important open problems in computational geometry is that of characterizing and recognizing visibility graphs.

For the case of planar graphs, we will use the following result by T. Nishizeki extensively:

Theorem 1.2.4 *(Nishizeki [97]): Any planar 2-connected graph G with $n \geq 14$ vertices and minimum vertex degree greater than or equal to 3 has a maximum matching with at least $\lfloor \frac{n+4}{3} \rfloor$ edges. When $n \leq 14$, the number of edges in such a matching is $\lfloor \frac{n}{2} \rfloor$.*

The proofs of these results are not of concern in this work; the interested reader can find them in many books on graph theory.

1.3 Exercises

Ex. 1.1 Prove that any tree with n vertices contains exactly $n-1$ edges.

Ex. 1.2 Prove Theorem 1.2.3.

Ex. 1.3 Prove that the graph $GT(P)$ obtained from a triangulation of a polygon P has at least two vertices of degree two.

Ex. 1.4 Prove that any polygon can be triangulated, and that any triangulation of a polygon with n vertices has exactly $n-2$ triangles.

Ex. 1.5 Let H be a tree such that the degree of all its vertices is at most 3. Show that there is a polygon P and a triangulation T of P such that the dual of $GT(P)$ is H. Moreover, show that there is a polygon P such that P has a unique triangulation T whose dual is H.

Ex. 1.6 Show that every convex polyhedron in R^3 can be decomposed into a set of convex tetrahedrons.

Ex. 1.7 Find a convex polyhedron with n vertices that can be decomposed into tetrahedrons in two different ways such that the first decomposition has a linear number of tetrahedrons, and the second partition has a quadratic number.