Chapter 1
The Classification Theorem: Informal Presentation

1.1 Introduction

Few things are as rewarding as finally stumbling upon the view of a breathtaking landscape at the turn of a path after a long hike. Similar experiences occur in mathematics, music, art, etc. When we first read about the classification of the compact surfaces, we sensed that if we prepared ourself for a long hike, we could probably enjoy the same kind of exhilaration.

The Problem

Define a suitable notion of *equivalence* of surfaces so that *a complete list of representatives, one in each equivalence class of surfaces, is produced*, each representative having a simple explicit description called a *normal form*. By a suitable notion of equivalence, we mean that two surfaces S_1 and S_2 are equivalent iff there is a "nice" bijection between them.

The *classification theorem* for compact surfaces says that, despite the fact that surfaces appear in many diverse forms, surfaces can be classified, which means that every compact surface is equivalent to exactly one representative surface, also called a surface in *normal form*. Furthermore, there exist various kinds of normal forms that are very concrete, for example, polyhedra obtained by gluing the sides of certain kinds of regular planar polygons. For this type of normal form, there is also a finite set of transformations with the property that every surface can be transformed into a normal form in a finite number of steps.

Of course, in order to make the above statements rigorous, one needs to define precisely

1. What is a surface.
2. What is a suitable notion of equivalence of surfaces.
3. What are normal forms of surfaces.

J. Gallier and D. Xu, *A Guide to the Classification Theorem for Compact Surfaces*,
Geometry and Computing 9, DOI 10.1007/978-3-642-34364-3_1,

This is what we aim to do in this book!

For the time being, let us just say that a surface is a topological space with the property that around every point, there is an open subset that is homeomorphic to an open disc in the plane (the interior of a circle).[1] We say that a surface is *locally Euclidean*. Informally, two surfaces X_1 and X_2 are equivalent if each one can be continuously deformed into the other. More precisely, this means that there is a *continuous bijection*, $f : X_1 \to X_2$, such that f^{-1} is also continuous (we say that f is a *homeomorphism*). So, by "nice" bijection we mean a homeomorphism, and two surfaces are considered to be equivalent if there is a homeomorphism between them.

The Solution

Every proof of the classification theorem for compact surfaces comprises two steps:

1. *A topological step*. This step consists in showing that every compact surface *can be triangulated*.
2. *A combinatorial step*. This step consists in showing that every triangulated surface can be converted to a normal form in a finite number of steps, using some (finite) set of transformations.

To clarify step 1, we have to explain what is a *triangulated surface*. Intuitively, a surface can be triangulated if it is homeomorphic to a space obtained by pasting triangles together along edges. A technical way to achieve this is to define the combinatorial notion of a two-dimensional complex, a formalization of a polyhedron with triangular faces. We will explain thoroughly the notion of triangulation in Chap. 3 (especially Sect. 3.2).

The fact that every surface can be triangulated was first proved by Radó in 1925 (Fig. 1.1). This proof is also presented in Ahlfors and Sario [1] (see Chap. I, Sect. 8).

The proof is fairly complicated and the intuition behind it is unclear. Other simpler and shorter proofs have been found and we will present in Appendix E a proof due to Carsten Thomassen [15] which we consider to be the most easily accessible (if not the shortest).

There are a number of ways of implementing the combinatorial step. Once one realizes that a triangulated surface can be cut open and laid flat on the plane, it is fairly intuitive that such a flattened surface can be brought to normal form, but the

[1]More rigorously, we also need to require a surface to be Hausdorff and second-countable; see Definition 2.3.

Fig. 1.1 Tibor Radó, 1895–1965

details are a bit tedious. We will give a complete proof in Chap. 6 and a preview of this process in Sect. 1.2.

It should also be said that distinct normal forms of surfaces can be distinguished by simple invariants:

(a) Their *orientability* (orientable or non-orientable).
(b) Their *Euler–Poincaré characteristic*, an integer that encodes the number of "holes" in the surface.

Actually, it is not easy to define precisely the notion of orientability of a surface and to prove rigorously that the Euler–Poincaré characteristic is a topological invariant, which means that it is preserved under homeomorphisms.

Intuitively, the notion of orientability can be explained as follows. Let A and B be two bugs on a surface assumed to be transparent. Pick any point p, assume that A stays at p and that B travels along any closed curve on the surface starting from p dragging along a coin. A memorizes the coin's face at the beginning of the path followed by B. When B comes back to p after traveling along the closed curve, two possibilities may occur:

1. A sees the same face of the coin that he memorized at the beginning of the trip.
2. A sees the other face of the coin.

If case 1 occurs for all closed curves on the surface, we say that it is *orientable*. This will be the case for a sphere or a torus. However, if case 2 occurs, then we say that the surface is *nonorientable*. This phenomenon can be observed for the surface known as the *Möbius strip*, see Fig. 1.2.

Orientability will be discussed rigorously in Sect. 4.5 and the Euler–Poincaré characteristic and its invariance in Chap. 5 (see especially Theorem 5.2).

In the words of Milnor himself, the classification theorem for compact surfaces is a formidable result. This result was first proved rigorously by Brahana [2] in 1921 but it had been stated in various forms as early as 1861 by Möbius [12], by Jordan [8] in 1866, by von Dyck [4] in 1888 and by Dehn and Heegaard [3] in 1907, so it was the culmination of the work of many (see Appendix D).

Indeed, a rigorous proof requires, among other things, a precise definition of a surface and of orientability, a precise notion of triangulation, and a precise way of

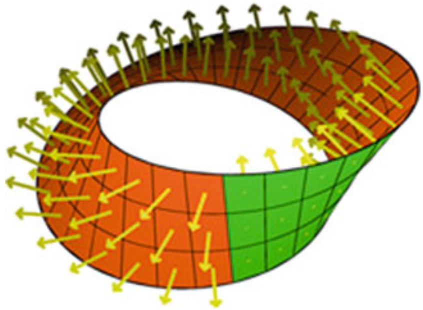

Fig. 1.2 A Möbius strip in $\mathbb{R}^3$ (K. Polthier of FU Berlin)

determining whether two surfaces are homeomorphic or not. This requires some notions of algebraic topology such as, fundamental groups, homology groups, and the Euler–Poincaré characteristic. Most steps of the proof are rather involved and it is easy to lose track.

One aspect of the proof that we find particularly fascinating is the use of certain kinds of graphs (called cell complexes) and of some kinds of rewrite rules on these graphs, to show that every triangulated surface is equivalent to some cell complex *in normal form*. This presents a challenge to researchers interested in rewriting, as the objects are unusual (neither terms nor graphs), and rewriting is really modulo cyclic permutations (in the case of boundaries). We hope that this book will inspire some of the researchers in the field of rewriting to investigate these mysterious rewriting systems.

Our goal is to help the reader reach the top of the mountain [the classification theorem for compact surfaces, with or without boundaries (also called borders)], and help him not to get lost or discouraged too early. This is not an easy task!

We provide quite a bit of topological background material and the basic facts of algebraic topology needed for understanding how the proof goes, with more than an impressionistic feeling.

We also review abelian groups and present a proof of the structure theorem for finitely generated abelian groups due to Pierre Samuel. Readers with a good mathematical background should proceed directly to Sect. 2.2, or even to Sect. 3.1.

We hope that this book will be helpful to readers interested in geometry, and who still believe in the rewards of serious hiking!

1.2 Informal Presentation of the Theorem

Until Riemann's work in the early 1850s, surfaces were always dealt with from a local point of view (as parametric surfaces) and topological issues were never considered. In fact, the view that a surface is a topological space locally homeomorphic to the Euclidean plane was only clearly articulated in the early 1930s by

Fig. 1.3 James W Alexander, 1888–1971 (*left*), Hassler Whitney, 1907–1989 (*middle*) and Herman K H Weyl, 1885–1955 (*right*)

Fig. 1.4 Bernhard Riemann, 1826–1866 (*left*), August Ferdinand Möbius, 1790–1868 (*middle left*), Johann Benedict Listing, 1808–1882 (*middle right*) and Camille Jordan, 1838–1922 (*right*)

Alexander and Whitney (although Weyl also adopted this view in his seminal work on Riemann surfaces as early as 1913) (Fig. 1.3).

After Riemann, various people, such as Listing, Möbius and Jordan, began to investigate topological properties of surfaces, in particular, *topological invariants* (Fig. 1.4). Among these invariants, they considered various notions of connectivity, such as the maximum number of (non self-intersecting) closed pairwise disjoint curves that can be drawn on a surface without disconnecting it and, the Euler–Poincaré characteristic. These mathematicians took the view that a (compact) surface is made of some elastic stretchable material and they took for granted the fact that every surface can be triangulated. Two surfaces S_1 and S_2 were considered *equivalent* if S_1 could be mapped onto S_2 by a continuous mapping "without tearing and duplication" and S_2 could be similarly be mapped onto S_1. This notion of equivalence is a precursor of the notion of a *homeomorphism* (not formulated precisely until the 1900s) that is, an invertible map, $f\colon S_1 \to S_2$, such that both f and its inverse, f^{-1}, are continuous.

Möbius and Jordan seem to be the first to realize that the main problem about the topology of (compact) surfaces is to find invariants (preferably numerical) to decide the equivalence of surfaces, that is, to decide whether two surfaces are homeomorphic or not.

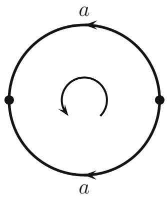

Fig. 1.5 A cell representing a sphere (boundary aa^{-1})

The crucial fact that makes the classification of compact surfaces possible is that every (connected) compact, triangulated surface can be opened up and laid flat onto the plane (as one connected piece) by making a finite number of cuts along well chosen simple closed curves on the surface.

Then, we may assume that the flattened surface consists of convex polygonal pieces, called *cells*, whose edges (possibly curved) are tagged with labels associated with the curves used to cut the surface open. Every labeled edge occurs twice, possibly shared by two cells.

Consequently, every compact surface can be obtained from a set of convex polygons (possibly with curved edges) in the plane, called cells, by gluing together pairs of unmatched edges.

These sets of cells representing surfaces are called *cell complexes*. In fact, it is even possible to choose the curves so that they all pass through a single common point and so, every compact surface is obtained from a single polygon with an even number of edges and whose vertices all correspond to a single point on the surface.

For example, a sphere can be opened up by making a cut along half of a great circle and then by pulling apart the two sides (the same way we open a Chinese lantern) and smoothly flattening the surface until it becomes a flat disk. Symbolically, we can represent the sphere as a round cell with two boundary curves labeled and oriented identically, to indicate that these two boundaries should be identified, see Fig. 1.5.

We can also represent the boundary of this cell as a string, in this case, aa^{-1}, by following the boundary counter-clockwise and putting an inverse sign on the label of an edge iff this edge is traversed in the opposite direction.

To open up a torus, we make two cuts: one using any half-plane containing the axis of revolution of the torus, the other one using a plane normal to the axis of revolution and tangential to the torus (see Fig. 1.6).

By deformation, we get a square with opposite edges labeled and oriented identically, see Fig. 1.7. The boundary of this square can be described by a string obtained by traversing it counter-clockwise: we get $aba^{-1}b^{-1}$, where the last two edges have an inverse sign indicating that they are traversed backwards.

A surface (orientable) with two holes can be opened up using four cuts. Observe that such a surface can be thought of as the result of gluing two tori together: take two tori, cut out a small round hole in each torus and glue them together along the boundaries of these small holes. Then, we make two cuts to split the two tori (using a plane containing the “axis” of the surface) and then two more cuts to open

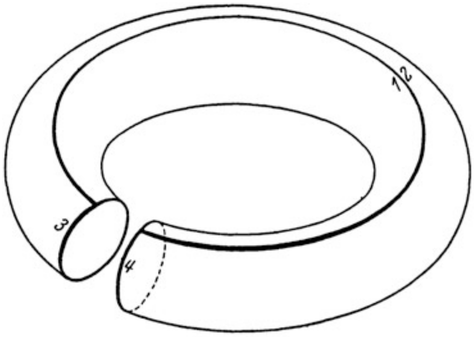

FIG. 284.

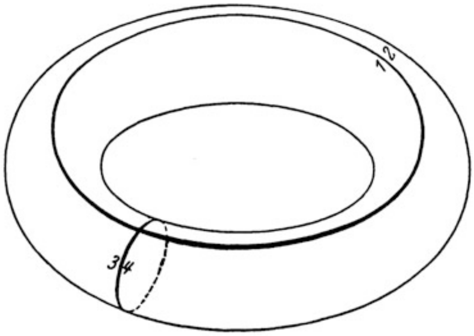

FIG. 285.

Fig. 1.6 Cutting open a torus, from Hilbert and Cohn–Vossen [6], p. 264

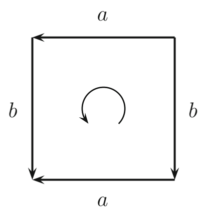

Fig. 1.7 A cell representing a torus (boundary $aba^{-1}b^{-1}$)

up the surface. This process is very nicely depicted in Hilbert and Cohn–Vossen [7] (pp. 300–301) and in Fréchet and Fan [5] (pp. 38–39), see Fig. 1.8.

The result is that a surface with two holes can be represented by an octagon with four pairs of matching edges, as shown in Fig. 1.9.

A surface (orientable) with three holes can be opened up using six cuts and is represented by a 12-gon with edges pairwise identified as shown in Cohn–Vossen [7] (pp. 300–301), see Fig. 1.8.

In general, an orientable surface with g holes (a surface of *genus* g) can be opened up using $2g$ cuts and can be represented by a regular $4g$-gon with edges pairwise identified, where the boundary of this $4g$-gon is of the form

$$a_1b_1a_1^{-1}b_1^{-1}a_2b_2a_2^{-1}b_2^{-1}\cdots a_gb_ga_g^{-1}b_g^{-1},$$

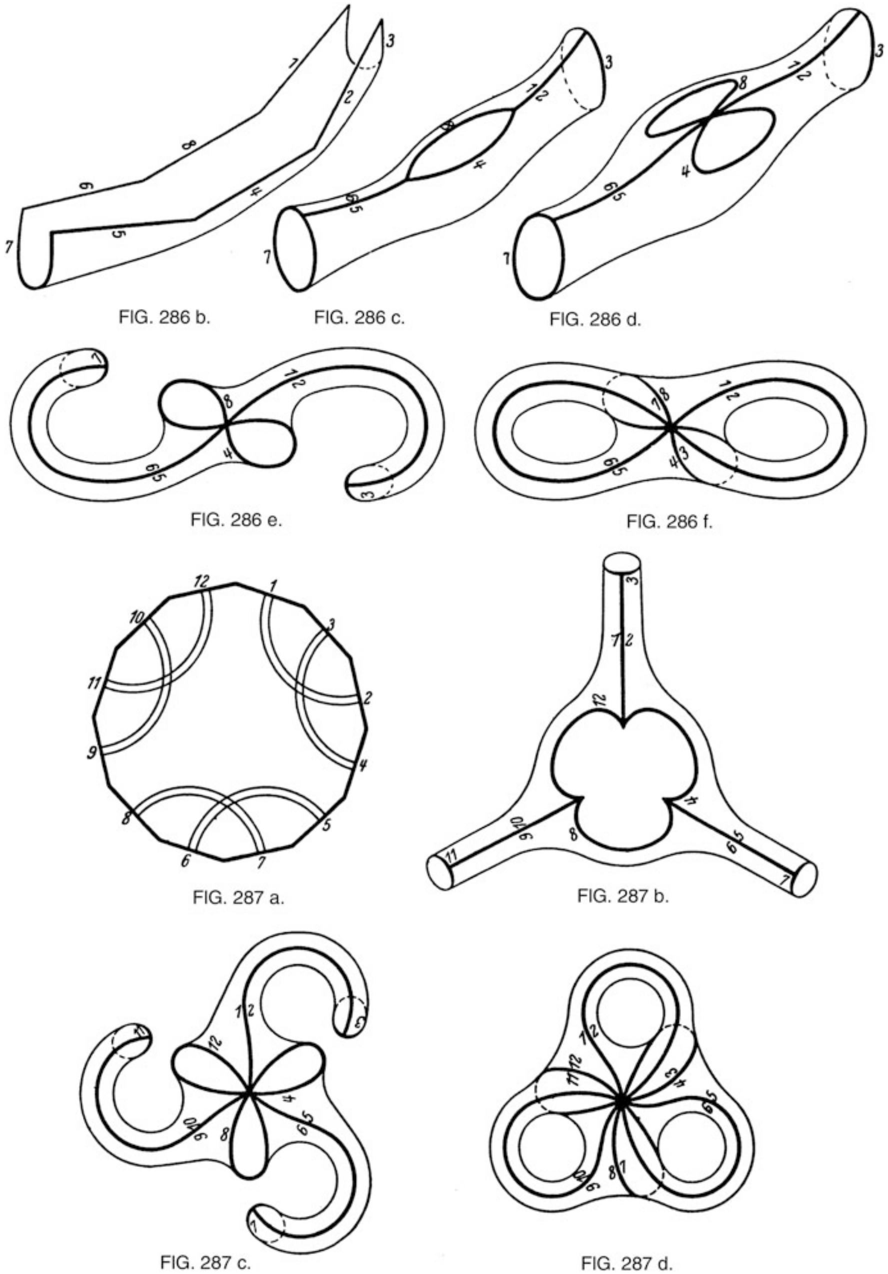

Fig. 1.8 Constructing a surface with two holes and a surface with three holes by gluing the edges of a polygon, from Hilbert and Cohn–Vossen [6], p. 265

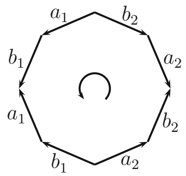

Fig. 1.9 A cell representing a surface with two holes (boundary $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$)

called type (I). The sphere is represented by a single cell with boundary

$$aa^{-1}, \text{ or } \epsilon \text{ (the empty string);}$$

this cell is also considered of type (I).

The normal form of type (I) has the following useful geometric interpretation: A torus can be obtained by gluing a “tube” (a bent cylinder) onto a sphere by cutting out two small disks on the surface of the sphere and then gluing the boundaries of the tube with the boundaries of the two holes. Therefore, we can think of a surface of type (I) as the result of attaching g handles onto a sphere. The cell complex, $aba^{-1}b^{-1}$, is called a *handle*.

In addition to being orientable or nonorientable, surfaces may have *boundaries*. For example, the first surface obtained by slicing a torus shown in Fig. 1.6 (FIG. 284) is a bent cylinder that has two boundary circles. Similarly, the top three surfaces shown in Fig. 1.8 (FIG. 286b–d) are surfaces with boundaries. On the other hand, the sphere and the torus have no boundary.

As we said earlier, every surface (with or without boundaries) can be triangulated, a fact proved by Radó in 1925. Then, the crucial step in proving the classification theorem for compact surfaces is to show that every triangulated surface can be converted to an equivalent one in *normal form*, namely, represented by a $4g$-gon in the orientable case or by a $2g$-gon in the nonorientable case, using some simple transformations involving cuts and gluing. This can indeed be done, and next we sketch the conversion to normal form for surfaces without boundaries, following a minor variation of the method presented in Seifert and Threlfall [14].

Since our surfaces are already triangulated, we may assume that they are given by a finite set of planar polygons with curved edges. Thus, we have a finite set, F, of faces, each face, $A \in F$, being assigned a boundary, $B(A)$, which can be viewed as a string of oriented edges from some finite set, E, of edges. In order to deal with oriented edges, we introduce the set, E^{-1}, of “inverse” edges and we assume that we have a function, $B\colon F \to (E \cup E^{-1})^*$, assigning a string or oriented edges, $B(A) = a_1 a_2 \cdots a_n$, to each face, $A \in F$, with $n \geq 2$.[2] Actually, we also introduce the set, F^{-1}, of inversely oriented faces A^{-1}, with the convention that $B(A^{-1}) = a_n^{-1} \cdots a_2^{-1} a_1^{-1}$ if $B(A) = a_1 a_2 \cdots a_n$. We also do not distinguish

[2] In Sect. 6.1, we will allow $n \geq 0$.

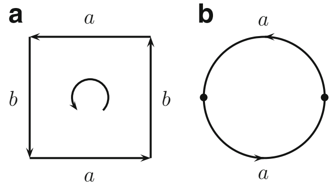

Fig. 1.10 (**a**) A projective plane (boundary $abab$). (**b**) A projective plane (boundary aa)

between boundaries obtained by cyclic permutations. We call A and A^{-1} *oriented faces*. Every finite set, K, of faces representing a surface satisfies two conditions:

(1) Every oriented edge, $a \in E \cup E^{-1}$, occurs twice as an element of a boundary. In particular, this means that if a occurs twice in some boundary, then it does not occur in any other boundary.
(2) K is connected. This means that K is not the union of two disjoint systems satisfying condition (1).

A finite (nonempty) set of faces with an assignment of boundaries satisfying conditions (1) and (2) is called a *cell complex*. We already saw examples of cell complexes at the beginning of this section. For example, a torus is represented by a single face with boundary $aba^{-1}b^{-1}$. A more precise definition of a cell complex will be given in Definition 6.1.

Every oriented edge has a source vertex and a target vertex, but distinct edges may share source or target vertices. Now this may come as a surprise, but the definition of a cell complex allows other surfaces besides the familiar ones, namely *nonorientable* surfaces. For example, if we consider a single cell with boundary $abab$, as shown in Fig. 1.10a, we have to construct a surface by gluing the two edges labeled a together, but this requires first "twisting" the square piece of material by an angle π, and similarly for the two edges labeled b.

One will quickly realize that there is no way to realize such a surface without self-intersection in $\mathbb{R}^3$ and this can indeed be proved rigorously although it is nontrivial; see Note F.1. The above surface is the *real projective plane*, $\mathbb{RP}^2$.

As a topological space, the real projective plane is the set of all lines through the origin in $\mathbb{R}^3$. A more concrete representation of this space is obtained by considering the upper hemisphere,

$$S^2_+ = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1,\ z \geq 0\}.$$

Now, every line through the origin not contained in the plane $z = 0$ intersects the upper hemisphere, S^2_+, in a single point, whereas every line through the origin contained in the plane $z = 0$ intersects the equatorial circle in two antipodal points. It follows that the projective plane, $\mathbb{RP}^2$, can be viewed as the upper hemisphere, S^2_+, with antipodal on its boundary identified. This is not easy to visualize! Furthermore, the orthogonal projection along the z-axis yields a bijection between S^2_+ and the closed disk,

$$\overline{D} = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\},$$

so the projective plane, $\mathbb{RP}^2$, can be viewed as the closed disk, $\overline{D}$, with antipodal on its boundary identified. This explains why the cell in Fig. 1.10a yields the projective plane by identification of edges and so does the circular cell with boundary aa shown in Fig. 1.10b. A way to realize the projective plane as a surface in $\mathbb{R}^3$ with self-intersection is shown in Note F.2. Other methods for realizing $\mathbb{RP}^2$ are given in Appendix A.

Let us go back to the notion of orientability. This is a subtle notion and coming up with a precise definition is harder than one might expect. The crucial idea is that if a surface is represented by a cell complex, then this surface is orientable if there is a way to assign a direction of traversal (clockwise or counterclockwise) to the boundary of every face, so that when we fold and paste the cell complex by gluing together every edge a with its inverse a^{-1}, no tearing or creasing takes place. The result of the folding and pasting process should be a surface in $\mathbb{R}^3$. In particular, the gluing process does not involve any twist and does not cause any self-intersection.

Another way to understand the notion of orientability is that if we start from some face A_0 and follow a closed path $A_0, A_1, \ldots, A_n$ on the surface by moving from each face A_i to the next face A_{i+1} if A_i and A_{i+1} share a common edge, then when we come back to $A_0 = A_n$, the orientation of A_0 has not changed. Here is a rigorous way to capture the notion of orientability.

Given a cell complex, K, an *orientation of* K is a set of faces $\{A^\epsilon \mid A \in F\}$, where each face A^ϵ is obtained by choosing one of the two oriented faces A, A^{-1} for every face $A \in F$, that is, $A^\epsilon = A$ or $A^\epsilon = A^{-1}$. An orientation is *coherent* if every edge a in $E \cup E^{-1}$ occurs once in the set of boundaries of the faces in $\{A^\epsilon \mid A \in F\}$. A cell complex, K, is *orientable* if is has some coherent orientation.

For example, the complex with boundary $aba^{-1}b^{-1}$ representing the torus is orientable, but the complex with boundary aa representing the projective plane is not orientable. The cell complex K with two faces A_1 and A_2 whose boundaries are given by $B(A_1) = abc$ and $B(A_2) = bac$ is orientable since we can pick the orientation $\{A_1, A_2^{-1}\}$. Indeed, $B(A_2^{-1}) = c^{-1}a^{-1}b^{-1}$ and every oriented edge occurs once in the faces in $\{A_1, A_2^{-1}\}$; see Fig. 1.11. Note that the orientation of A_2 is the opposite of the orientation shown on the Figure, which is the orientation of A_1.

It is clear that every surface represented by a normal form of type (I) is orientable. It turns out that every nonorientable surface (with $g \geq 1$ "holes") can be represented by a $2g$-gon where the boundary of this $2g$-gon is of the form

$$a_1a_1a_2a_2 \cdots a_ga_g,$$

called type (II). All these facts will be proved in Chap. 6, Sect. 6.3.

The normal form of type (II) also has a useful geometric interpretation: Instead of gluing g handles onto a sphere, glue g projective planes, *i.e.* cross-caps, onto a sphere. The cell complex with boundary, aa, is called a *cross-cap* (Fig. 1.10(b)).

Another famous nonorientable surface known as the *Klein bottle* is obtained by gluing matching edges of the cell shown in Fig. 1.13a. This surface was first described by Klein [9] (1882) (Fig. 1.12). As for the projective plane, using the results of Note F.1, it can be shown that the Klein bottle cannot be embedded in $\mathbb{R}^3$.

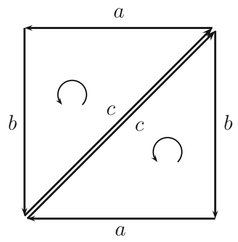

Fig. 1.11 An orientable cell complex with $B(A_1) = abc$ and $B(A_2) = bac$

Fig. 1.12 Felix C Klein, 1849–1925

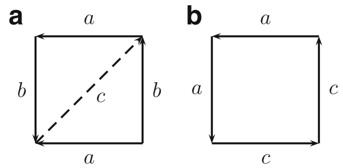

Fig. 1.13 (**a**) A Klein bottle (boundary $aba^{-1}b$). (**b**) A Klein bottle (boundary $aacc$)

If we cut the cell shown in Fig. 1.13a along the edge labeled c and then glue the resulting two cells (with boundaries abc and $bc^{-1}a^{-1}$) along the edge labeled b, we get the cell complex with boundary $aacc$ showed in Fig. 1.13b. Therefore, the Klein bottle is the result of gluing together two projective planes by cutting out small disks in these projective planes and then gluing them along the boundaries of these disks. However, in order to obtain a representation of a Klein bottle in $\mathbb{R}^3$ as a surface with a self-intersection it is better to use the edge identification specified by the cell complex of Fig. 1.13a. First, glue the edges labeled a together, obtaining a tube (a cylinder), then twist and bend this tube to let it penetrate itself in order to glue the edges labeled b together, see Fig. 1.14. Other pictures of a Klein bottle are shown in Fig. 1.15.

In summary, there are two kinds *normal forms* of cell complexes: these cell complexes $K = (F, E, B)$ in normal form have a single face A ($F = \{A\}$), and either

(I) $E = \{a_1, \ldots, a_p, b_1, \ldots, b_p\}$ and

$$B(A) = a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_p b_p a_p^{-1} b_p^{-1},$$

where $p \geq 0$, or

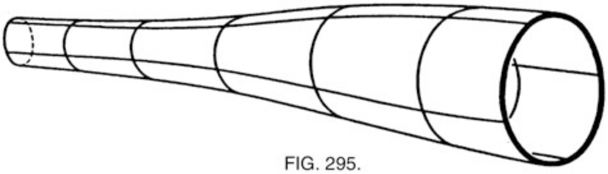

FIG. 295.

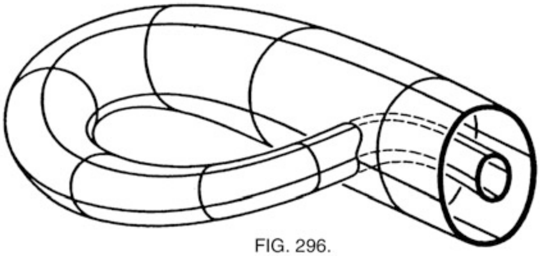

FIG. 296.

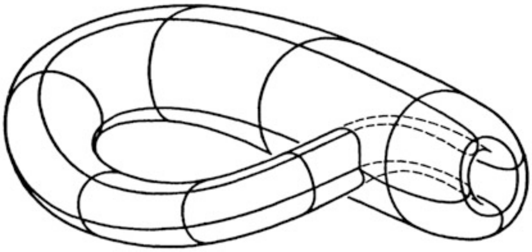

FIG. 297.

Fig. 1.14 Construction of a Klein bottle, from Hilbert and Cohn–Vossen [6], pp. 271–272

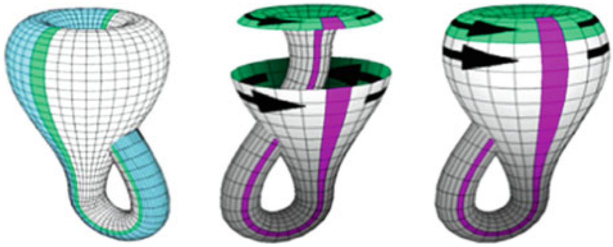

Fig. 1.15 Klein bottles in $\mathbb{R}^3$ (K. Polthier of FU Berlin)

(II) $E = \{a_1, \ldots, a_p\}$ and

$$B(A) = a_1 a_1 \cdots a_p a_p,$$

where $p \geq 1$.

Observe that canonical complexes of type (I) are orientable, whereas canonical complexes of type (II) are not. When $p = 0$, the canonical complex of type (I)

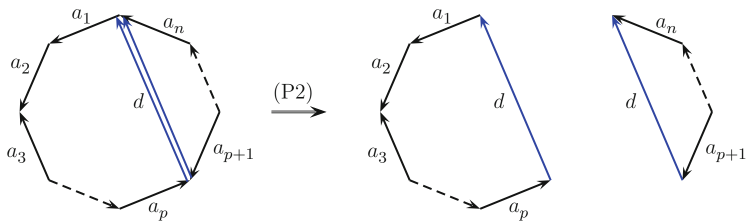

Fig. 1.16 Rule (P2)

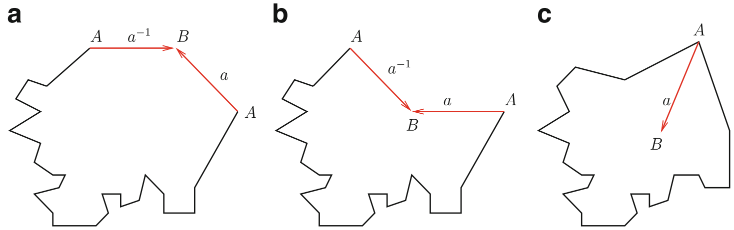

Fig. 1.17 Elimination of aa^{-1}

corresponds to a sphere, and we let $B(A) = \epsilon$ (the empty string). The above surfaces have no boundary; the general case of surfaces with boundaries is covered in Chap. 6. Then, the combinatorial form the classification theorem for (compact) surfaces can be stated as follows:

Theorem 1.1. *Every cell complex K can be converted to a cell complex in normal form by using a sequence of steps involving a transformation $(P2)$ and its inverse: splitting a cell complex, and gluing two cell complexes together.*

Actually, to be more precise, we should also have an edge-splitting and an edge-merging operation but, following Massey [11], if we define the elimination of pairs aa^{-1} in a special manner, only one operation is needed, namely:

Transformation P2: Given a cell complex, K, we obtain the cell complex, K', by *elementary subdivision of* K (or *cut*) if the following operation, (P2), is applied: Some face A in K with boundary $a_1 \ldots a_p a_{p+1} \ldots a_n$ is replaced by two faces A' and A'' of K', with boundaries $a_1 \ldots a_p d$ and $d^{-1} a_{p+1} \ldots a_n$, where d is an edge in K' and not in K. Of course, the corresponding replacement is applied to A^{-1}.

Rule (P2) is illustrated in Fig. 1.16.

Proof (Sketch of proof for Theorem 1.1). The procedure for converting a cell complex to normal form consists of several steps.

Step 1. Elimination of strings aa^{-1} in boundaries, see Fig. 1.17.

Step 2. Vertex Reduction.

The purpose of this step is to obtain a cell complex with a single vertex. We first perform step 1 repeatedly until all occurrences of the form aa^{-1} have been

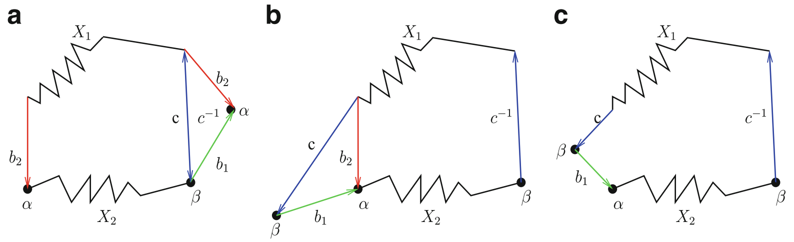

Fig. 1.18 Reduction to a single inner vertex

eliminated. If the remaining sequence has no edges left, then it must be of type (I).

Otherwise, consider an inner vertex $\alpha = (b_1, \ldots, b_m)$. If α is not the only inner vertex, then there is another inner vertex β. We assume without loss of generality that b_1 is the edge that connects β to α. Also, we must have $m \geq 2$, since otherwise there would be a string $b_1 b_1^{-1}$ in some boundary. Thus, locate the string $b_1 b_2^{-1}$ in some boundary. Suppose it is of the form $b_1 b_2^{-1} X_1$, and using (P2), we can split it into $b_1 b_2^{-1} c$ and $c^{-1} X_1$ (see Fig. 1.18a). Now locate b_2 in the boundary, suppose it is of the form $b_2 X_2$. Since b_2 differs from b_1, b_1^{-1}, c, c^{-1}, we can eliminate b_2 by applying $(\mathrm{P2})^{-1}$. This is equivalent to cutting the triangle $c b_1 b_2^{-1}$ off along edge c, and pasting it back with b_2 identified with b_2^{-1} (see Fig. 1.18b).

This has the effect of shrinking α. Indeed, as one can see from Fig. 1.18c, there is one less vertex labeled α, and one more labeled β.

This procedure can be repeated until $\alpha = (b_1)$, at which stage b_1 is eliminated using step 1. Thus, it is possible to eliminate all inner vertices except one. Thus, from now on, we will assume that there is a single inner vertex.

Step 3. Reduction to a single face and introduction of cross-caps.

We may still have several faces. We claim that for every face A, if there is some face B such that $B \neq A$, $B \neq A^{-1}$, and there is some edge a both in the boundary of A and in the boundary of B, due to the fact that all faces share the same inner vertex, and thus all faces share at least one edge. Thus, if there are at least two faces, from the above claim and using $(\mathrm{P2})^{-1}$, we can reduce the number of faces down to one. It it easy to check that no new vertices are introduced.

Next, if some boundary contains two occurrences of the same edge a, i.e., it is of the form $aXaY$, where X, Y denote strings of edges, with $X, Y \neq \epsilon$, we show how to make the two occurrences of a adjacent. This is the attempt to group the cross-caps together, resulting in a sequence that denotes a cell complex of type (II).

The above procedure is essentially the same as the one we performed in our vertex reduction step. The only difference is that we are now interested in the edge sequence in the boundary, not the vertices. The rule shows that by

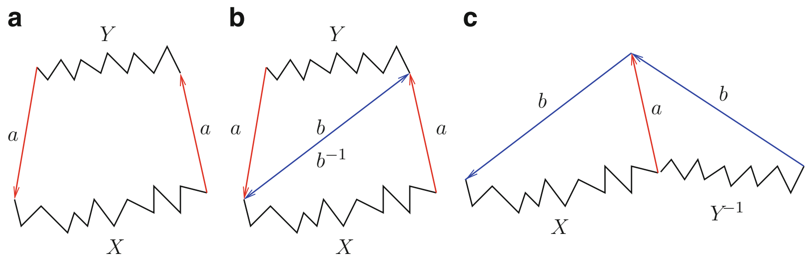

Fig. 1.19 Grouping the cross-caps

introducing a new edge b and its inverse, we can cut the cell complex in two along the new edge, and then paste the two parts back by identifying the the two occurrences of the same edge a, resulting in a new boundary with a cross-cap, as shown in Fig. 1.19c. By repeating step 3, we convert boundaries of the form $aXaY$ to boundaries with cross-caps.

Step 4. Introduction of handles.

The purpose of this step is to convert boundaries of the form $aUbVa^{-1}Xb^{-1}Y$ to boundaries $cdc^{-1}d^{-1}YXVU$ containing handles. This is the attempt to group the handles together, resulting in a sequence that denotes a cell complex of type (I). See Fig. 1.20.

Each time the rewrite rule is applied to the boundary sequence, we introduce a new edge and its inverse to the polygon, and then cut and paste the same way as we have described so far. Iteration of this step preserves cross-caps and handles.

Step 5. Transformation of handles into cross-caps.

At this point, one of the last obstacles to the canonical form is that we may still have a mixture of handles and cross-caps. If a boundary contains a handle and a cross-cap, the trick is to convert a handle into two cross-caps. This can be done in a number of ways. Massey [11] shows how to do this using the fact that the connected sum of a torus and a Möbius strip is equivalent to the connected sum of a Klein bottle and a Möbius strip. We prefer to explain how to convert a handle into two cross-caps using four applications of the cut and paste method using rule (P2) and its inverse, as presented in Seifert and Threlfall [14] (Sect. 38).

The first phase is to split a cell as shown in Fig. 1.21a into two cells using a cut along a new edge labeled d and then two glue the resulting new faces along the two edges labeled c, obtaining the cell showed in Fig. 1.21b. The second phase is to split the cell in Fig. 1.21b using a cut along a new edge labeled a_1 and then glue the resulting new faces along the two edges labeled b, obtaining the cell showed in Fig. 1.21c. The third phase is to split the cell in Fig. 1.22c using a cut along a new edge labeled a_2 and then glue the resulting new faces along the two edges labeled a, obtaining the cell showed in Fig. 1.22d. Finally, we split the cell in Fig. 1.22d using a cut along a new edge labeled a_3 and then glue the resulting new faces along the two edges labeled d, obtaining the cell showed in Fig. 1.22e.

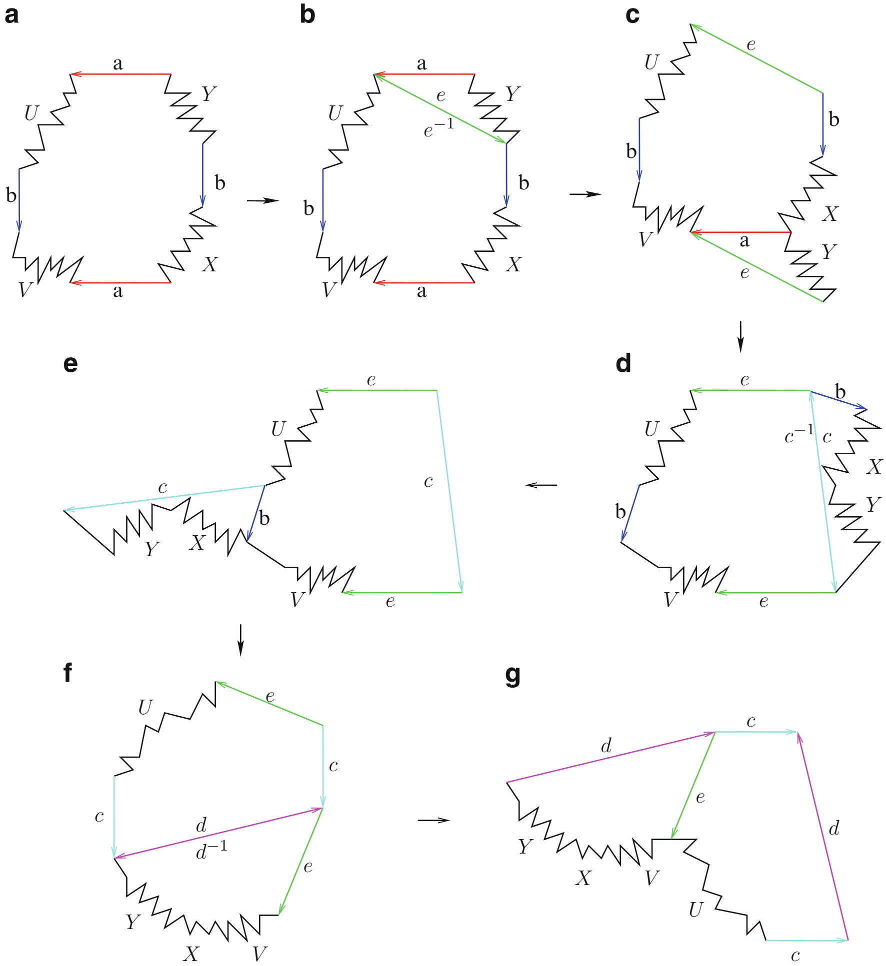

Fig. 1.20 Grouping the handles

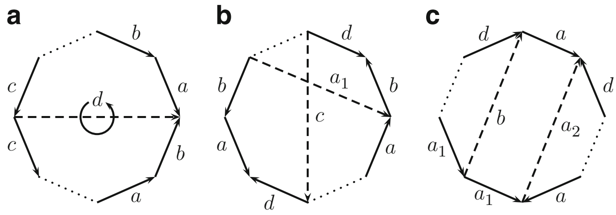

Fig. 1.21 Step 5, phases 1 and 2

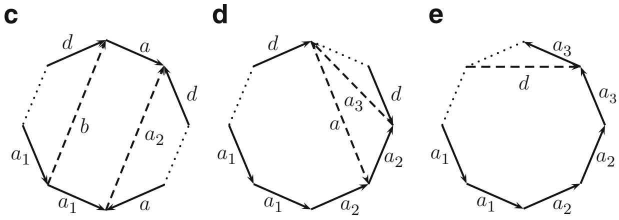

Fig. 1.22 Step 5, phases 3 and 4

Note that in the cell showed in Fig. 1.22e, the handle $aba^{-1}b^{-1}$ and the cross-cap cc have been replaced by the three consecutive cross-caps, $a_1a_1a_2a_2a_3a_3$.

Using the above procedure, every compact surface represented as a cell complex can be reduced to normal form, which proves Theorem 1.1. □

The next step is to show that distinct normal forms correspond to inequivalent surfaces, that is, surfaces that are not homeomorphic.

First, it can be shown that the orientability of a surface is preserved by the transformations for reducing to normal form. Second, if two surfaces are homeomorphic, then they have the same nature of orientability. The difficulty in this step is to define properly what it means for a surface to be orientable; this is done in Sect. 4.5 using the degree of a map in the plane.

Third, we can assign a numerical invariant to every surface, its *Euler–Poincaré characteristic*. For a triangulated surface K, if n_0 is the number of vertices, n_1 is the number of edges, and n_2 is the number of triangles, then the Euler–Poincaré characteristic of K is defined by

$$\chi(K) = n_0 - n_1 + n_2.$$

Then, we can show that homeomorphic surfaces have the same Euler–Poincaré characteristic and that distinct normal forms with the same type of orientability have different Euler–Poincaré characteristics. It follows that any two distinct normal forms correspond to inequivalent surfaces. We obtain the following version of the classification theorem for compact surfaces:

Theorem 1.2. *Two compact surfaces are homeomorphic iff they agree in character of orientability and Euler–Poincaré characteristic.*

Actually, Theorem 1.2 is a special case of a more general theorem applying to surfaces with boundaries as well (Theorem 6.2). All this will be proved rigorously in Chap. 6. Proving rigorously that the Euler–Poincaré characteristic is a topological invariant of surfaces will require a fair amount of work. In fact, we will have to define homology groups. In any case, we hope that the informal description of the reduction to normal form given in this section has raised our reader's curiosity enough to entice him to read the more technical development that follows.

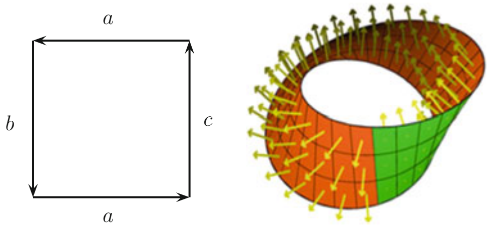

Fig. 1.23 *Left*: A Möbius strip (boundary $abac$). *Right*: A Möbius strip in $\mathbb{R}^3$ (K. Polthier of FU Berlin)

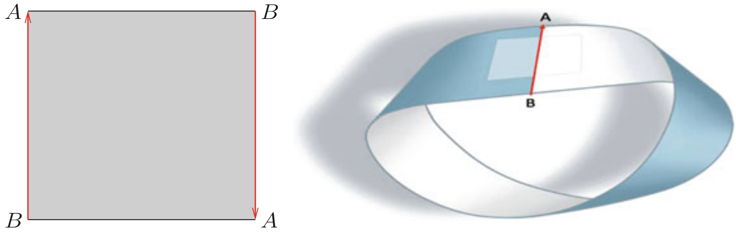

Fig. 1.24 Construction of a Möbius strip

To close this introductory chapter, let us go back briefly to surfaces with boundaries. Then, there is a well-known nonorientable surface realizable in $\mathbb{R}^3$, the *Möbius strip*. This surface was discovered independently by Listing [10] (1862) and Möbius [13] (1865).

The Möbius strip is obtained from the cell complex in Fig. 1.23 by gluing the two edges labeled a together. Observe that this requires a twist by π in order to glue the two edges labeled a properly.

The resulting surface shown in Fig. 1.23 and in Fig. 1.24 has a single boundary since the two edges b and c become glued together, unlike the situation where we do not make a twist when gluing the two edges labeled a, in which case we get a torus with two distinct boundaries, b and c.

It turns out that if we cut out a small hole into a projective plane we get a Möbius strip. This fact is nicely explained in Fréchet and Fan [5] (p. 42) or Hilbert and Cohn–Vossen [7] (pp. 315–316). It follows that we get a realization of a Möbius band with a flat boundary if we remove a small disk from a cross-cap. For this reason, this version of the Möbius strip is often called a cross-cap. Furthermore, the Klein bottle is obtained by gluing two Möbius strips along their boundaries (see Fig. 1.25). This is shown in Massey [11] using the cut and paste method, see Chap. 1, Lemma 7.1.

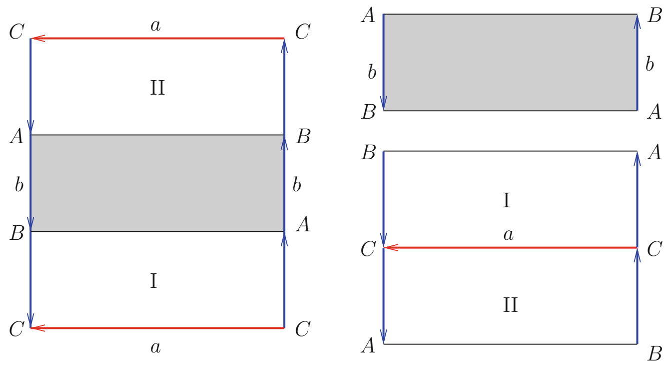

Fig. 1.25 Construction of a Klein bottle from two Möbius strips

References

1. L.V. Ahlfors, L. Sario, *Riemann Surfaces*, Princeton Math. Series, vol. 2 (Princeton University Press, Princeton, 1960)
2. H.R. Brahana, Systems of circuits on two-dimensional manifolds. Ann. Math. **23**(2), 144–168 (1921)
3. M. Dehn, P. Heegaard, in *III AB.3 Analysis Situs*, ed. by W.F. Meyer. Encyklopädie Der Mathematischen Wissenschaften Mit Einschluss Ihrer Anwendungen (Druck Und Verlag Von B.G. Teubner, Leipzig, 1907), pp. 153–220
4. W. von Dyck, Beiträge zur analysis situs. Mathematische Annalen **32**, 457–512 (1888)
5. M. Fréchet, K. Fan, *Invitation to Combinatorial Topology*, 1st edn. (Dover, New York, 2003)
6. D. Hilbert, S. Cohn–Vossen, *Anschauliche Geometrie*, 2nd edn. (Springer, New York, 1996)
7. D. Hilbert, S Cohn–Vossen, *Geometry and the Imagination* (Chelsea, New York, 1952)
8. C. Jordan, Sur la déformation des surfaces. J. de Mathématiques Pures et Appliquées 2^e série **11**, 105–109 (1866)
9. F. Klein, *Über Riemanns Theorie der Algebraischen Funktionen und Ihrer Integrale*, 1st edn. (Druck Und Verlag Von B.G. Teubner, Leipzig, 1882)
10. J.B. Listing, Der census räumlicher complexe oder verallgemeinerung des euler'schen satzes von den polyëdern. Abhandlungen der Königlichen Gesellschaft der Wissenschaften in Göttingen **10**, 97–182 (1862)
11. W.S. Massey, *A Basic Course in Algebraic Topology*, Graduate Texts in Mathematics, vol. 127, 1st edn. (Springer, New York, 1991)
12. A.F. Möbius, Zur theorie der polyëder und der elementarverwandtschaft. Oeuvres Complètes, Tome **2**, 519–559 (1861)
13. A.F. Möbius, Ueber die bestimmung des inhaltes eines polyëders. Berichte über die Verhandlungen der Königlich Sächsischen Gesellschaft der Wissenschaften, Mathematisch-physikalische Klasse **17**, 31–68 (1865)
14. H. Seifert, W. Threlfall, *A Textbook of Topology*, 1st edn. (Academic, Boston, 1980)
15. C. Thomassen, The Jordan-Schonflies Theorem and the classification of surfaces. Am. Math. Monthly **99**(2), 116–131 (1992)

http://www.springer.com/978-3-642-34363-6

A Guide to the Classification Theorem for Compact Surfaces
Gallier, J.; Xu, D.
2013, XII, 178 p. 78 illus., 20 illus. in color., Hardcover
ISBN: 978-3-642-34363-6