The Classification Theorem for Compact Surfaces

Jean Gallier

CIS Department
University of Pennsylvania

jean@cis.upenn.edu

October 31, 2012

The Classification of Surfaces: The Problem

So you think you know what a surface is?

The Classification of Surfaces: The Problem

So you think you know what a surface is?

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*.

The Classification of Surfaces: The Problem

So you think you know what a surface is?

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*.

The *classification theorem* says that, despite the fact that surfaces appear in many diverse forms, every compact surface is equivalent to exactly one representative surface, also called a surface in *normal form*.

The Problem

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*.

The Problem

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*.

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.

The Problem

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*.

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.

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.

Abstract Surfaces

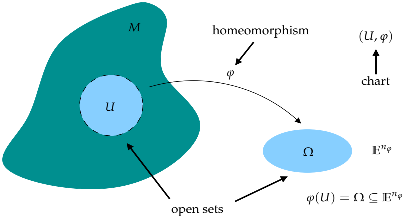

Figure : An abstract surface

Abstract Surfaces

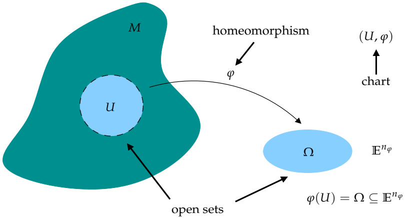

Figure : An abstract surface

The surface is covered by a finite number of open subsets.

Abstract Surfaces

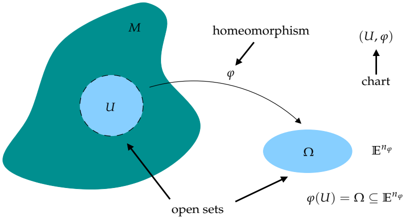

Figure : An abstract surface

The surface is covered by a finite number of open subsets.

Two surfaces X_1 and X_2 are equivalent if each one can be *continuously deformed* into the other. This means that there is a *continuous bijection*, $f\colon X_1 \to X_2$, such that f^{-1} is also continuous (a *homeomorphism*).

The Classification of Surfaces: The Solution

Every proof of the classification theorem comprises two steps:

The Classification of Surfaces: The Solution

Every proof of the classification theorem 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.

The Classification of Surfaces: The Solution

Every proof of the classification theorem 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*.

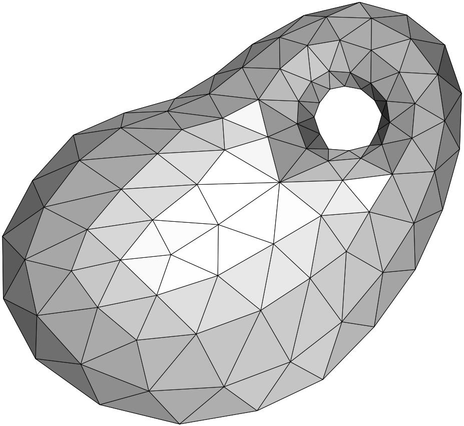

Figure : A triangulated surface

The Solution

A technical way to achieve this is to define the combinatorial notion of a *2-dimensional complex*, a formalization of a polyhedron with triangular faces.

The Solution

A technical way to achieve this is to define the combinatorial notion of a *2-dimensional complex*, a formalization of a polyhedron with triangular faces.

Every surface can be triangulated; first proved by Radó in 1925.

Figure : Tibor Radó, 1895-1965

The Solution

The proof is fairly complicated. The intuition behind it is unclear. The most easily accessible proof (if not the shortest) is due to Carsten Thomassen and uses *graph theory*.

The Solution

The proof is fairly complicated. The intuition behind it is unclear. The most easily accessible proof (if not the shortest) is due to Carsten Thomassen and uses *graph theory*.

There are a number of ways of implementing the combinatorial step.

The Solution

The proof is fairly complicated. The intuition behind it is unclear. The most easily accessible proof (if not the shortest) is due to Carsten Thomassen and uses *graph theory*.

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. The details are tedious.

The Solution

Distinct normal forms of surfaces can be distinguished by simple invariants:

The Solution

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.

The Solution

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.

Not so easy to define precisely the notion of orientability, and to prove rigorously that the Euler–Poincaré characteristic is a topological invariant.

Figure : Dog Logic

Orientability

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 begining of the path followed by B.

Orientability

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 begining of the path followed by B.

When B comes back to p after traveling along the closed curve, two possibilites 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.

Orientability

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 begining of the path followed by B.

When B comes back to p after traveling along the closed curve, two possibilites 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 is the case for a sphere or a torus.

Orientability

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*.

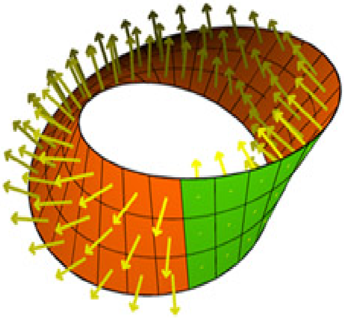

Figure : A Möbius strip in $\mathbb{R}^3$ (Image courtesy of Konrad Polthier of FU Berlin)

Tools Needed for the Proof

This result was first proved rigorously by Brahana in 1921, but it had been stated in various forms as early as 1861 by Möbius, by Jordan in 1866, by von Dyck in 1888 and by Dehn and Heegaard in 1907.

Tools Needed for the Proof

This result was first proved rigorously by Brahana in 1921, but it had been stated in various forms as early as 1861 by Möbius, by Jordan in 1866, by von Dyck in 1888 and by Dehn and Heegaard in 1907.

A rigorous proof requires: a precise definition of a surface and of orientability, a precise notion of triangulation, and a precise way of determining whether two surfaces are homeomorphic or not.

Tools Needed for the Proof

This result was first proved rigorously by Brahana in 1921, but it had been stated in various forms as early as 1861 by Möbius, by Jordan in 1866, by von Dyck in 1888 and by Dehn and Heegaard in 1907.

A rigorous proof requires: a precise definition of a surface and of orientability, a precise notion of triangulation, and a precise way of 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.

Main Idea of the Proof

Until Riemann's work in the early 1850's, surfaces were always dealt with from a *local* point of view (as parametric surfaces) and topological issues were never considered.

Main Idea of the Proof

Until Riemann's work in the early 1850's, surfaces were always dealt with from a *local* point of view (as parametric surfaces) and topological issues were never considered.

The view that a surface is a topological space locally homeomorphic to the Euclidean plane was only clearly articulated in the early 1930's by Alexander and Whitney (although Weyl also adopted this view in his seminal work on Riemann surfaces as early as 1913).

Main Idea of the Proof

Until Riemann's work in the early 1850's, surfaces were always dealt with from a *local* point of view (as parametric surfaces) and topological issues were never considered.

The view that a surface is a topological space locally homeomorphic to the Euclidean plane was only clearly articulated in the early 1930's by Alexander and Whitney (although Weyl also adopted this view in his seminal work on Riemann surfaces as early as 1913).

After Riemann, various people, such as Listing, Möbius and Jordan, began to investigate topological properties of surfaces, in particular, *topological invariants*.

Figure : James W Alexander, 1888- 1971 (left), Hassler Whitney, 1907-1989 (middle) and Herman K H Weyl, 1885-1955 (right)

Invariants: 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.

Invariants: 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 strechable material and they took for granted the fact that every surface can be triangulated.

Invariants: 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 strechable 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.

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

Figure : 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)

Crucial Fact

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.

Crucial Fact

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.

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.

Crucial Fact

These sets of cells representing surfaces are called *cell complexes*.

Crucial Fact

These sets of cells representing surfaces are called *cell complexes*.

It is even possible to choose the curves so that they all pass through a single common point.

Crucial Fact

These sets of cells representing surfaces are called *cell complexes*.

It is even possible to choose the curves so that they all pass through a single common point.

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.

Crucial Fact

These sets of cells representing surfaces are called *cell complexes*.

It is even possible to choose the curves so that they all pass through a single common point.

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.

A sphere can be opened up by making a cut along half of a great circle and then by pulling apart the two sides and smoothly flattening the surface until it becomes a flat disk (Chinese lantern).

Symbolically: the sphere is a round cell with two boundary curves labeled and oriented identically, to indicate that these two boundaries should be identified.

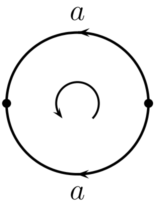

Figure : A cell representing a sphere (boundary aa^{-1})

To open up a torus, we make two cuts.

To open up a torus, we make two cuts.

By deformation, we get a square with opposite edges labeled and oriented identically:

To open up a torus, we make two cuts.

By deformation, we get a square with opposite edges labeled and oriented identically:

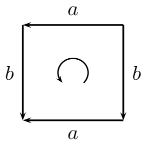

Figure : A cell representing a torus (boundary $aba^{-1}b^{-1}$)

To open up a torus, we make two cuts.

By deformation, we get a square with opposite edges labeled and oriented identically:

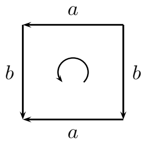

Figure : A cell representing a torus (boundary $aba^{-1}b^{-1}$)

The boundary can be described by the string $aba^{-1}b^{-1}$.

fallen. Die übrigen Seiten 3, 4, die dabei in Kreise übergegangen sind, können wir in der vorgeschriebenen Weise aneinanderheften, indem wir (Abb. 284) den Kreiszylinder verbiegen. Wir erhalten schließlich die Fläche des Torus, und der Rand der Rechtecksfläche geht in ein „kanonisches Schnittsystem" des Torus über, wobei jede Kurve zwei Randstrecken der Rechtecksfläche entspricht (Abb. 285 und 275b). Umgekehrt: schneidet man den Torus längs eines kanonischen Systems auf, so ergibt sich stets eine Figur, die dem Rechteck mit der angegebenen Ränderzuordnung topologisch äquivalent ist. Man kann dieses Verfahren auf alle „Brezeln" ausdehnen. Hat die Brezel den Zusammenhang $2p+1$, so besteht das kanonische Schnittsystem aus $2p$ Kurven. Die Zerschneidung liefert also ein $4p$-Eck mit einer bestimmten Ränderzuordnung. Für die Fälle $h=5, 7$, also $p=2, 3$, ist die Konstruktion durch Abb. 286, 287 veranschaulicht.

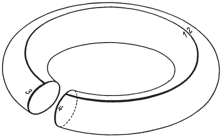

Abb. 284.

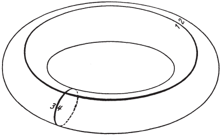

Abb. 285.

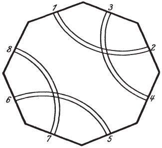

Abb. 286 a.

Die Abbildung der Brezeln auf $4p$-Ecke spielt sowohl in der Theorie der stetigen Abbildungen (vgl. S. 284) als auch in der Funktionentheorie (vgl. S. 294) eine wichtige Rolle. In beiden Anwendungen geht man davon aus, daß die $4p$-Ecke eine reguläre Gebietseinteilung der hyperbolischen Ebene (bzw. für $p=1$ der euklidischen Ebene) liefern, wie wir das auf S. 228 erörtert haben.

Wenn man die Ränderzuordnung abändert, erhält man außer den Brezeln noch eine große Anzahl weiterer Flächen, von denen uns einige im folgenden beschäftigen werden.

A surface (orientable) with two holes can be opened up using four cuts.

A surface (orientable) with two holes can be opened up using four cuts.

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.

Make two cuts to split the two tori (using a plane containing the "axis" of the surface), and then two more cuts to open up the surface.

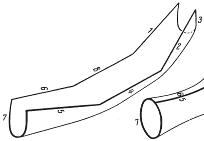

Abb. 286 b.

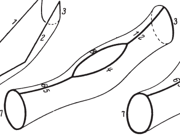

Abb. 286 c.

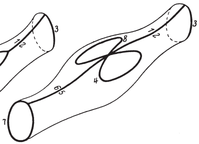

Abb. 286 d.

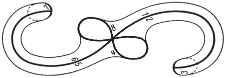

Abb. 286 e.

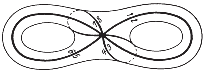

Abb. 286 f.

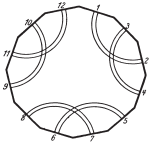

Abb. 287 a.

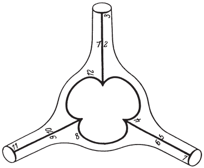

Abb. 287 b.

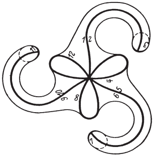

Abb. 287 c.

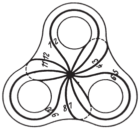

Abb. 287 d.

A surface with two holes can be represented by an octagon with four pairs of matching edges:

A surface with two holes can be represented by an octagon with four pairs of matching edges:

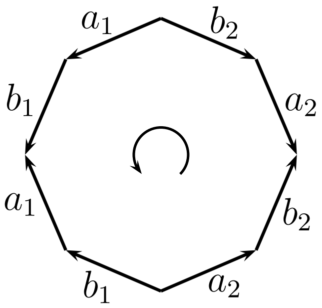

Figure : A cell representing a surface with two holes (boundary $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$)

A surface with two holes can be represented by an octagon with four pairs of matching edges:

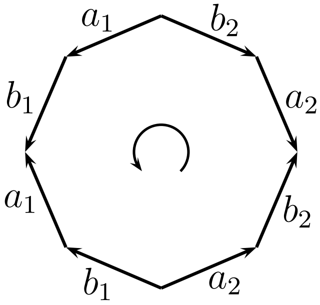

Figure : A cell representing a surface with two holes (boundary $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$)

A surface (orientable) with three holes can be opened up using 6 cuts and is represented by a 12-gon with edges pairwise identified.

Normal Form of Type I

An orientable surface with g holes (a surface of *genus* g) can be opened up using $2g$ cuts. It is represented by a regular $4g$-gon with edges pairwise identified. The boundary of this $4g$-gon is of the form

$$a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1},$$

called type (I).

Normal Form of Type I

An orientable surface with g holes (a surface of *genus* g) can be opened up using $2g$ cuts. It is represented by a regular $4g$-gon with edges pairwise identified. The boundary of this $4g$-gon is of the form

$$a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1},$$

called type (I).

The sphere is represented by a single cell with boundary

$$aa^{-1} (\text{or } \epsilon);$$

this cell is also considered of type (I).

We can think of a surface of type (I) as the result of attaching g handles onto a sphere.

We can think of a surface of type (I) as the result of attaching g handles onto a sphere.

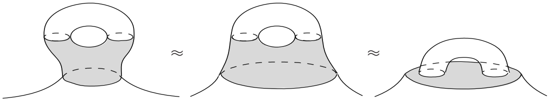

Fig. 6.7: Connected sum with a torus versus attaching a handle.

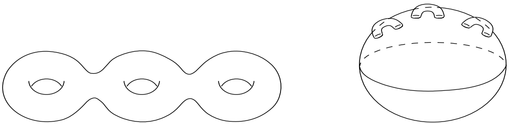

Figure : Attaching handles

We can think of a surface of type (I) as the result of attaching g handles onto a sphere.

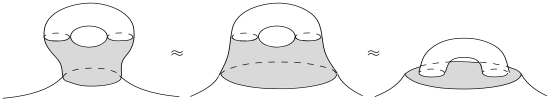

Fig. 6.7: Connected sum with a torus versus attaching a handle.

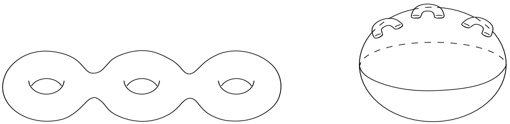

Figure : Attaching handles

The cell complex, $aba^{-1}b^{-1}$, is called a *handle*.

Cell Complexes

Since our surfaces are already triangulated, we may assume that they are given by a finite set of planar polygons with curved edges.

Cell Complexes

Since our surfaces are already triangulated, we may assume that they are given by a finite set of planar polygons with curved edges.

Cell complex = a finite set, F, of *faces*, each face, $A \in F$, being assigned a *boundary*, $B(A)$, a *string of oriented edges* from some finite set, E, of edges.

Cell Complexes

Since our surfaces are already triangulated, we may assume that they are given by a finite set of planar polygons with curved edges.

Cell complex = a finite set, F, of *faces*, each face, $A \in F$, being assigned a *boundary*, $B(A)$, a *string of oriented edges* from some finite set, E, of edges.

To deal with oriented edges: the set, E^{-1}, of "inverse" edges.

Cell Complexes

Since our surfaces are already triangulated, we may assume that they are given by a finite set of planar polygons with curved edges.

Cell complex = a finite set, F, of *faces*, each face, $A \in F$, being assigned a *boundary*, $B(A)$, a *string of oriented edges* from some finite set, E, of edges.

To deal with oriented edges: the set, E^{-1}, of "inverse" edges.

$B \colon F \to (E \cup E^{-1})^*$, assigns a string or oriented edges, $B(A) = a_1 a_2 \cdots a_n$, to each face, $A \in F$, with $n \geq 2$.

Cell Complexes

Also need the set, F^{-1}, of inversely oriented faces A^{-1}. Convention: $B(A^{-1}) = a_n^{-1} \cdots a_2^{-1} a_1^{-1}$ if $B(A) = a_1 a_2 \cdots a_n$. We do not distinguish between boundaries obtained by cyclic permutations.

Cell Complexes

Also need the set, F^{-1}, of inversely oriented faces A^{-1}. Convention: $B(A^{-1}) = a_n^{-1} \cdots a_2^{-1} a_1^{-1}$ if $B(A) = a_1 a_2 \cdots a_n$. We do not distinguish between boundaries obtained by cyclic permutations.

We call A and A^{-1} *oriented faces*.

Cell Complexes

Also need the set, F^{-1}, of inversely oriented faces A^{-1}. Convention: $B(A^{-1}) = a_n^{-1} \cdots a_2^{-1} a_1^{-1}$ if $B(A) = a_1 a_2 \cdots a_n$. We do not distinguish 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.

(2) K is connected.

Cell Complexes

A finite (nonempty) set of faces with an assignment of boundaries satisfying conditions (1) and (2) is called a *cell complex*.

Cell Complexes

A finite (nonempty) set of faces with an assignment of boundaries satisfying conditions (1) and (2) is called a *cell complex*.

For example, a torus is represented by a single face with boundary $aba^{-1}b^{-1}$.

Cell Complexes

A finite (nonempty) set of faces with an assignment of boundaries satisfying conditions (1) and (2) is called a *cell complex*.

For example, a torus is represented by a single face with boundary $aba^{-1}b^{-1}$.

Surprise: the definition of a cell complex allows other surfaces besides the familiar ones: *nonorientable* surfaces.

Nonorientable Surfaces

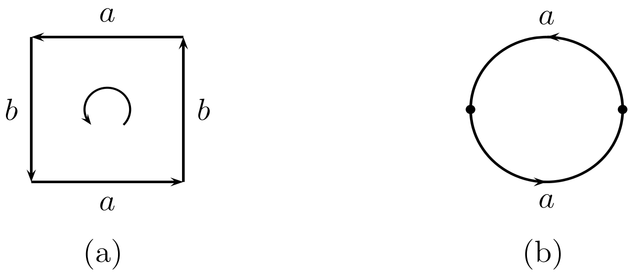

Figure : (a) A projective plane (boundary *abab*). (b) A projective plane (boundary *aa*).

We have to glue the two edges labeled *a* together. This requires first "twisting" the square piece of material by an angle π, and similarly for the two edges labeled *b*.

There is no way to realize such a surface without self-intersection in $\mathbb{R}^3$.

There is no way to realize such a surface without self-intersection in $\mathbb{R}^3$.

The above surface is the *real projective plane*, $\mathbb{RP}^2$.

There is no way to realize such a surface without self-intersection in $\mathbb{R}^3$.

The above surface is the *real projective plane*, $\mathbb{RP}^2$.

$\mathbb{RP}^2$ can be realized in $\mathbb{R}^3$ as the *cross-cap*.

projektive Ebene stets in dieser Weise in vier Dreiecke zerlegen[1]. In der EULERschen Formel hat man jetzt $E = 3$, $K = 6$, $F = 4$ zu setzen und erhält wieder $h = 2$.

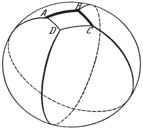

Abb. 303.

Wie wir die Ringfläche und die KLEINsche Fläche aus ihren Quadratmodellen durch Zusammenheftung erhalten haben, so wollen wir auch mit dem Quadratmodell der projektiven Ebene verfahren. Zu diesem Zweck verzerren wir das Quadrat zunächst in die Gestalt einer Kugelfläche, aus der ein kleines Viereck $ABCD$ herausgeschnitten ist (Abb. 303). Wir haben AB mit CD und DA mit BC zusammenzuheften. Das wird möglich, wenn wir die Fläche in den Punkten A und C heben und bei B und D senken und A und C sowie B und D einander nähern (Abb. 304). Schließlich erhalten wir eine geschlossene Fläche mit einer Strecke als Durchdringungslinie (Abb. 305). Sie ist topologisch der projektiven Ebene äquivalent.

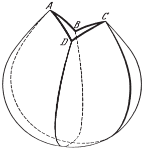

Abb. 304.

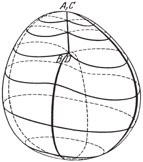

Abb. 305.

Es gibt eine algebraische Fläche, die diese Gestalt besitzt (Abb. 306). Ihre Gleichung ist

$$(k_1 x^2 + k_2 y^2)(x^2 + y^2 + z^2) - 2z(x^2 + y^2) = 0.$$

Die Fläche steht im Zusammenhang mit einer differentialgeometrischen Konstruktion. Wir gehen aus von einem Punkt P auf irgendeiner Fläche F, die in P konvex gekrümmt ist. Für alle Normalschnitte dieser Fläche in P (vgl. S. 162, 163) konstruieren wir die Krümmungskreise in P. Diese Kreisschar überstreicht dann gerade die in Abb. 306 dargestellte Fläche, ihre Durchdringungsstrecke ist ein Stück der in P errichteten Normalen der Ausgangsfläche; die angeführte Gleichung bezieht sich

[1] Die in Abb. 301 und 302 gezeichneten Einteilungen der projektiven Ebene hatten wir S. 131, 132 durch Projektion des Oktaeders erhalten.

Orientability

Orientability: a subtle notion.

Orientability

Orientability: a subtle notion.

Given a cell complex, K, an *orientation of K* is obtained by choosing one of the two oriented faces A, A^{-1} for every face $A \in F$; denote it A^{ϵ}.

Orientability

Orientability: a subtle notion.

Given a cell complex, K, an *orientation of K* is obtained by choosing one of the two oriented faces A, A^{-1} for every face $A \in F$; denote it A^{ϵ}.

An orientation is *coherent* if every edge a in $E \cup E^{-1}$ *occurs at most once* in the set of boundaries of the faces in $\{A^{\epsilon} \mid A \in F\}$.

Orientability

Orientability: a subtle notion.

Given a cell complex, K, an *orientation of K* is obtained by choosing one of the two oriented faces A, A^{-1} for every face $A \in F$; denote it A^{ϵ}.

An orientation is *coherent* if every edge a in $E \cup E^{-1}$ *occurs at most 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.

Orientability

Orientability: a subtle notion.

Given a cell complex, K, an *orientation of K* is obtained by choosing one of the two oriented faces A, A^{-1} for every face $A \in F$; denote it A^{ϵ}.

An orientation is *coherent* if every edge a in $E \cup E^{-1}$ *occurs at most 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.

The complex with boundary $aba^{-1}b^{-1}$ (a torus) is orientable, but the complex with boundary aa (the projective plane) is not orientable.

Normal Form of Type II

Every surface with normal form of type (I) is orientable.

Normal Form of Type II

Every surface with normal form of type (I) is orientable.

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

$$a_1 a_1 a_2 a_2 \cdots a_g a_g,$$

called type (II).

Normal Form of Type II

Every surface with normal form of type (I) is orientable.

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

$$a_1 a_1 a_2 a_2 \cdots a_g a_g,$$

called type (II).

Normal form of type (II): glue g projective planes, *i.e.* cross-caps (boundary aa), onto a sphere.

Normal Forms

In summary: two kinds *normal forms* (or *canonical forms*). These cell complexes $K = (F, E, B)$ in normal form have a single face A ($F = \{A\}$), and either

(I) $E = \{a_1, \ldots, a_g, b_1, \ldots, b_g\}$ and

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

where $g \geq 0$, or

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

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

where $g \geq 1$.

Canonical complexes of type (I) are orientable, whereas canonical complexes of type (II) are not.

Canonical complexes of type (I) are orientable, whereas canonical complexes of type (II) are not.

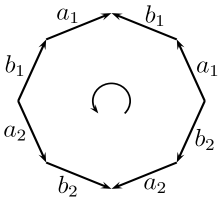

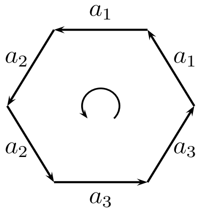

Figure : Examples of Normal Forms: (a) Type I; (b) Type II.

Euler–Poincaré characteristic

Given a cell complex, $K = (F, E, B)$, let $n_0 =$ the number of *vertices*, $n_1 =$ the number of elements in E (*edges*), and $n_2 =$ the number of elements in F (*faces*).

Euler–Poincaré characteristic

Given a cell complex, $K = (F, E, B)$, let n_0 = the number of *vertices*, n_1 = the number of elements in E (*edges*), and n_2 = the number of elements in F (*faces*).

The *Euler–Poincaré characteristic* of K is

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

Figure : Leonhard Euler, 1707-1783 (left) and Henri Poincaré, 1854-1912 (right)

A complex in normal form has a single vertex and a single face.

A complex in normal form has a single vertex and a single face.

Normal form of type I, $2g$ edges,

$$\chi(K) = 2 - 2g.$$

Normal form of type II, g edges,

$$\chi(K) = 2 - g.$$

A complex in normal form has a single vertex and a single face.

Normal form of type I, $2g$ edges,

$$\chi(K) = 2 - 2g.$$

Normal form of type II, g edges,

$$\chi(K) = 2 - g.$$

1. $\chi(\text{sphere}) = 2$ (since $g = 0$; $\epsilon \equiv aa^{-1}$),
2. $\chi(\text{torus}) = 0$ (since $g = 1$; $aba^{-1}b^{-1}$),
3. $\chi(2\text{−hole torus}) = -2$ (since $g = 2$; $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$),
4. $\chi(\text{projective plane}) = 1$ (since $g = 1$; aa),
5. $\chi(\text{Klein bottle}) = 0$ (since $g = 2$; $aabb$).

Cell Complexes and Triangulations

Every cell complex K defines a topological space $|K|$ obtained by a quotient process (identification of edges). Not clear that $|K|$ is surface.

Cell Complexes and Triangulations

Every cell complex K defines a topological space $|K|$ obtained by a quotient process (identification of edges). Not clear that $|K|$ is surface.

We can prove that $|K|$ is a surface by showing that every cell complex can be refined to a triangulated 2-complex.

Cell Complexes and Triangulations

Every cell complex K defines a topological space $|K|$ obtained by a quotient process (identification of edges). Not clear that $|K|$ is surface.

We can prove that $|K|$ is a surface by showing that every cell complex can be refined to a triangulated 2-complex.

A *triangulated 2-complex* K is a 2-complex satisfying some extra conditions so that $|K|$ is a surface.

Theorem 1

Every cell complex K can be refined to a triangulated 2-complex.

Theorem 1

Every cell complex K can be refined to a triangulated 2-complex.

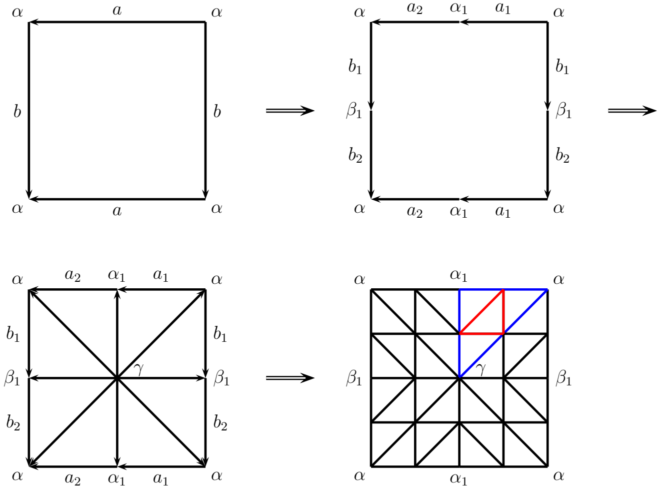

Theorem 1 implies that for every cell complex K, the space $|K|$ is a compact surface.

Theorem 1 implies that for every cell complex K, the space $|K|$ is a compact surface.

Theorem 2

Given any two finite triangulated 2-complexes K_1 and K_2, if $|K_1|$ and $|K_2|$ are homeomorphic, then K_1 and K_2 have the same character of orientability and $\chi(K_1) = \chi(K_2)$.

Distinct canonical complexes are inequivalent

The fact that $\chi(K_1) = \chi(K_2)$ follows form the fact that *homeomorphic spaces have isomorphic homology groups*, and that for a finite simplicial complex K of dimension m,

$$\chi(K) = \sum_{p=0}^{m} (-1)^p n_p = \sum_{p=0}^{m} (-1)^p \, r(H_p(K)),$$

where n_p is the number of p-simplicies in K.

Distinct canonical complexes are inequivalent

The fact that $\chi(K_1) = \chi(K_2)$ follows form the fact that *homeomorphic spaces have isomorphic homology groups*, and that for a finite simplicial complex K of dimension m,

$$\chi(K) = \sum_{p=0}^{m} (-1)^p n_p = \sum_{p=0}^{m} (-1)^p \, r(H_p(K)),$$

where n_p is the number of p-simplicies in K.

Consequence: any two distinct canonical complexes are not homeomorphic.

Combinatorial form the classification theorem

Theorem 3

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.

Transformation P2

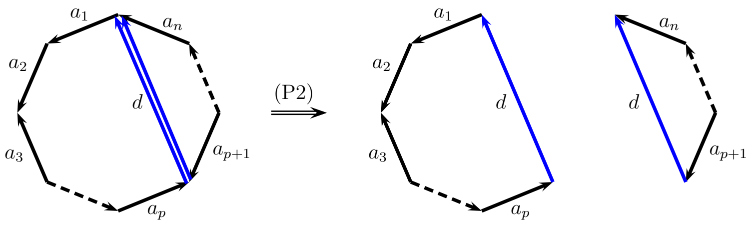

Figure : Transformation (P2)

First, (P2) is applied to the one-cell with boundary $aba^{-1}b$ to obtain a cell complex with two faces with boundaries, abc and $c^{-1}a^{-1}b$.

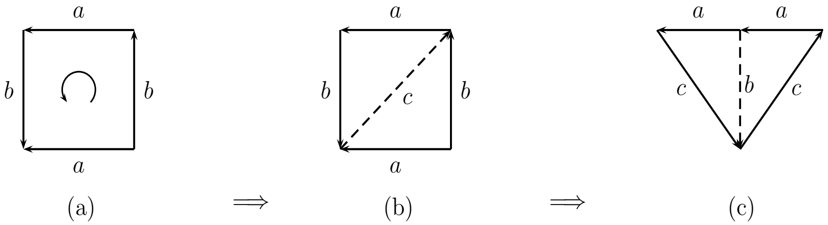

Figure : Example of elementary subdivision (P2) and its inverse

First, (P2) is applied to the one-cell with boundary $aba^{-1}b$ to obtain a cell complex with two faces with boundaries, abc and $c^{-1}a^{-1}b$.

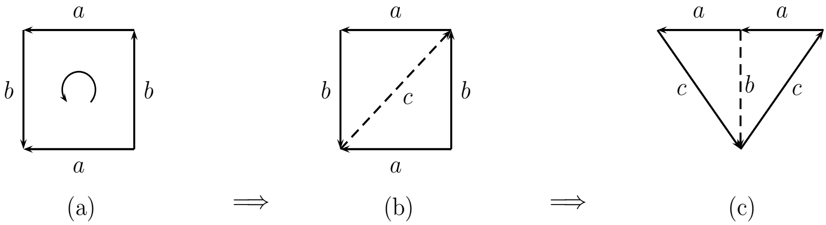

Figure : Example of elementary subdivision (P2) and its inverse

Then, these two faces are glued along the edge labeled b using $(\mathrm{P2})^{-1}$. Get a complex with boundary $aacc$: a *Klein bottle*.

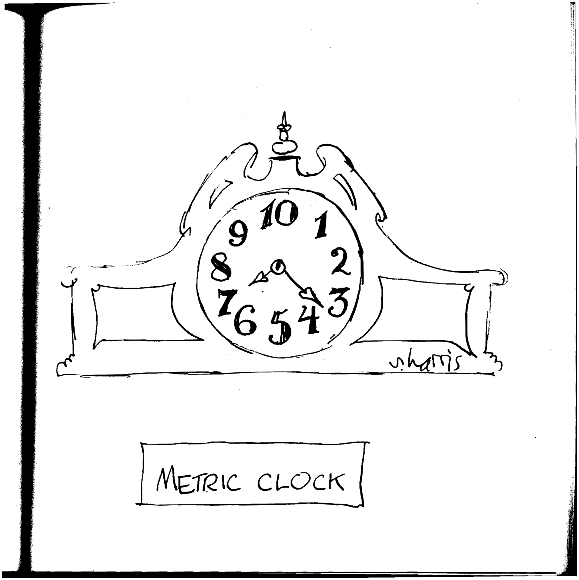

Figure : Metric Clock

The Final Theorem

Theorem 4

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

The Final Theorem

Theorem 4

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

There is a generalization of Theorem 4 to surfaces with boundaries.

Universal Covering Spaces of Surfaces

We can tile the plane with copies of the fundamental rectangle shown in blue (torus, boundary $aba^{-1}b^{-1}$) :

Universal Covering Spaces of Surfaces

We can tile the plane with copies of the fundamental rectangle shown in blue (torus, boundary $aba^{-1}b^{-1}$) :

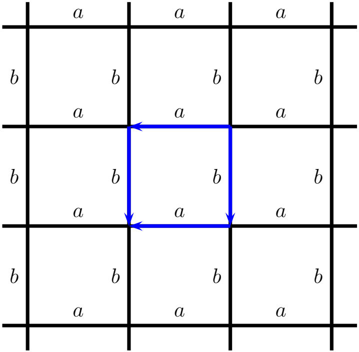

Figure : Tiling of the plane with tori

Every other tile is obtained by translation. This set of translations forms a group G generated by two translations (translate left one square, translate up one square).

Every other tile is obtained by translation. This set of translations forms a group G generated by two translations (translate left one square, translate up one square).

The torus T^2 = quotient of $\mathbb{R}^2$ under the *action of the group* G; quotient map

$$\pi \colon \mathbb{R}^2 \to T^2$$

makes the plane $\mathbb{R}^2$ a *universal cover* of the torus (as a smooth 2-manifold).

Every other tile is obtained by translation. This set of translations forms a group G generated by two translations (translate left one square, translate up one square).

The torus T^2 = quotient of $\mathbb{R}^2$ under the *action of the group* G; quotient map

$$\pi\colon \mathbb{R}^2 \to T^2$$

makes the plane $\mathbb{R}^2$ a *universal cover* of the torus (as a smooth 2-manifold).

The map π can be used to transfer the metric of $\mathbb{R}^2$ onto the torus (as a flat torus).

Can we do something similar with the 2-hole torus?

The 2-hole torus is defined by the cell (8 vertices) whose boundary is $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$.

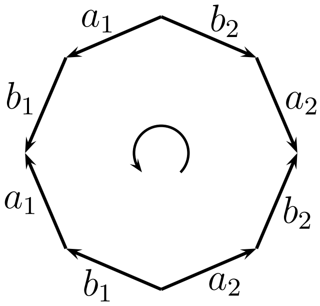

Figure : A cell representing a surface with two holes

Can we do something similar with the 2-hole torus?

The 2-hole torus is defined by the cell (8 vertices) whose boundary is $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$.

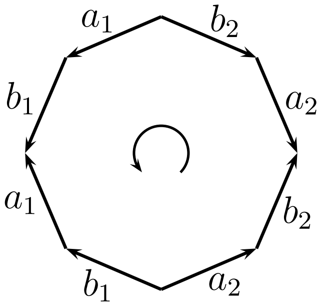

Figure : A cell representing a surface with two holes

Unfortunately, it is impossible to tile the plane with octagons!

If we allow curved edges, we can do it! Move to a *hyperbolic space*, such as the Poincaré disc (the interior of the unit disc).

If we allow curved edges, we can do it! Move to a *hyperbolic space*, such as the Poincaré disc (the interior of the unit disc).

In such a space, lines are arc of circles orthogonal to the boundary of the unit circle. Reflections become *inversions*.

If we allow curved edges, we can do it! Move to a *hyperbolic space*, such as the Poincaré disc (the interior of the unit disc).

In such a space, lines are arc of circles orthogonal to the boundary of the unit circle. Reflections become *inversions*.

It is possible to find a hyperbolic octogon whose angles are $\pi/4$, and to use inversions to fit other octagons to tile the hyperbolic plane.

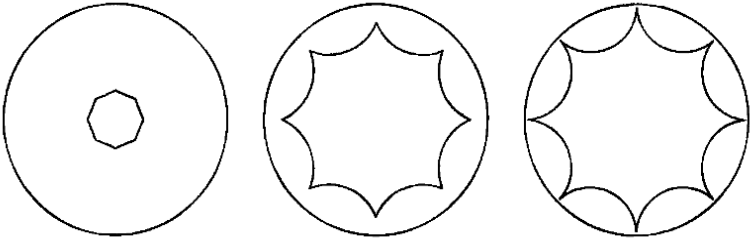

(a) (b) (c)

Figure 1.12. Bigger octagons in hyperbolic space have smaller angles. Between a tiny, Euclidean-like octagon with large angles (a) and a very large one with arbitrarily small angles (c) there must be one with angles exactly $\pi/4$ (b).

Figure : Finding a hyperbolic octagon (from Thurston)

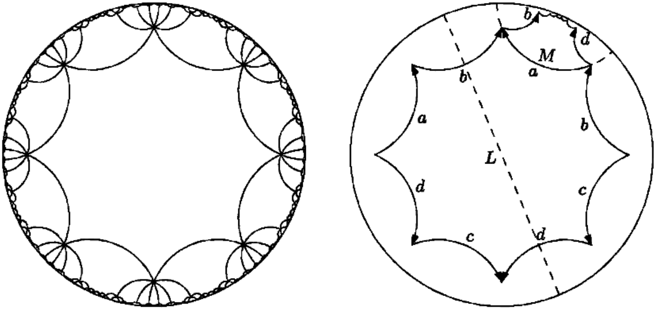

(a) (b)

Figure 1.13. A tiling of the hyperbolic plane by regular octagons. (a) A tiling of the hyperbolic plane by identical regular octagons, seen in the Poincaré disk projection. (b) To get the small octagon from the big one, reflect in L, then in M.

Figure : Tiling with hyperbolic octagons (from Thurston)

There is a group G of inversions that maps the fundamental domain to any other tile. The 2-hole torus $S =$ quotient of the hyperbolic plane $\mathbf{H}^2$ under the *action of the group* G.

There is a group G of inversions that maps the fundamental domain to any other tile. The 2-hole torus $S =$ quotient of the hyperbolic plane $\mathbf{H}^2$ under the *action of the group* G.

The quotient map

$$\pi\colon \mathbf{H}^2 \to S$$

makes $\mathbf{H}^2$ a *universal cover* of S (as a smooth 2-manifold). The metric of the hyperbolic space (a space of negative curvature) can be transferred to the surface S.

There is a group G of inversions that maps the fundamental domain to any other tile. The 2-hole torus $S =$ quotient of the hyperbolic plane $\mathbf{H}^2$ under the *action of the group* G.

The quotient map

$$\pi \colon \mathbf{H}^2 \to S$$

makes $\mathbf{H}^2$ a *universal cover* of S (as a smooth 2-manifold). The metric of the hyperbolic space (a space of negative curvature) can be transferred to the surface S.

The above process can be generalized to any smooth surface S of genus $g \geq 3$, so that the hyperbolic plane $\mathbf{H}^2$ is a universal cover of S.

There is a group G of inversions that maps the fundamental domain to any other tile. The 2-hole torus $S =$ quotient of the hyperbolic plane $\mathbf{H}^2$ under the *action of the group* G.

The quotient map

$$\pi \colon \mathbf{H}^2 \to S$$

makes $\mathbf{H}^2$ a *universal cover* of S (as a smooth 2-manifold). The metric of the hyperbolic space (a space of negative curvature) can be transferred to the surface S.

The above process can be generalized to any smooth surface S of genus $g \geq 3$, so that the hyperbolic plane $\mathbf{H}^2$ is a universal cover of S.

The torus and the Klein bottle have $\mathbb{R}^2$ as universal cover. The sphere S^2 is the universal cover of $\mathbb{RP}^2$.

There is a group G of inversions that maps the fundamental domain to any other tile. The 2-hole torus $S =$ quotient of the hyperbolic plane $\mathbf{H}^2$ under the *action of the group* G.

The quotient map

$$\pi\colon \mathbf{H}^2 \to S$$

makes $\mathbf{H}^2$ a *universal cover* of S (as a smooth 2-manifold). The metric of the hyperbolic space (a space of negative curvature) can be transferred to the surface S.

The above process can be generalized to any smooth surface S of genus $g \geq 3$, so that the hyperbolic plane $\mathbf{H}^2$ is a universal cover of S.

The torus and the Klein bottle have $\mathbb{R}^2$ as universal cover. The sphere S^2 is the universal cover of $\mathbb{RP}^2$.

This provides a kind of *global parametrization* for any smooth surface S.

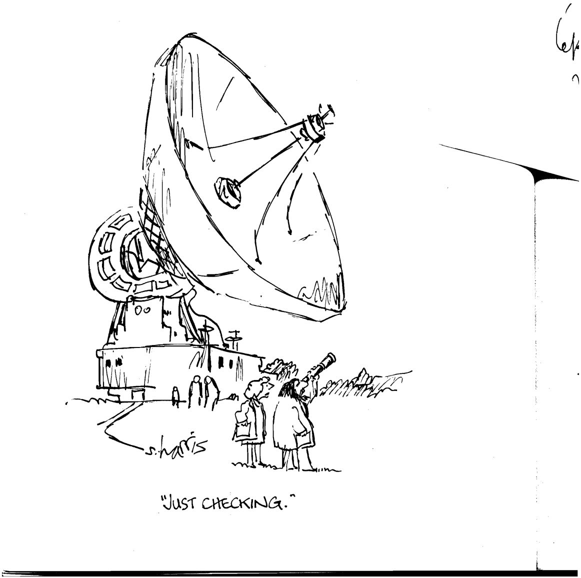

Figure : Just Checking

Some References

1. *Algebraic Topology: An Introduction*, Massey, William S., Springer Verlag, 1987, GTM No. 56.
2. *Invitation to Combinatorial Topology*, Fréchet, Maurice and Fan, Ky, Dover, 2003.
3. *A Textbook of Topology*, Seifert, H. and Threlfall, W., Academic Press, 1980 (originaly 1934).
4. *Riemann Surfaces*, Ahlfors, Lars V. and Sario, Leo, Princeton University Press, 1960.
5. *Three-Dimensional Topology and Geometry*, Volume 1, William Thurston, edited by Silvio Levy, Princeton University Press, 1997.
6. *A Guide to the Classification Theorem for Compact Surfaces*, Jean Gallier and Dianna Xu, Springer Verlag, to appear, 2013, http://www.cis.upenn.edu/∼jean/gbooks/surftop.html

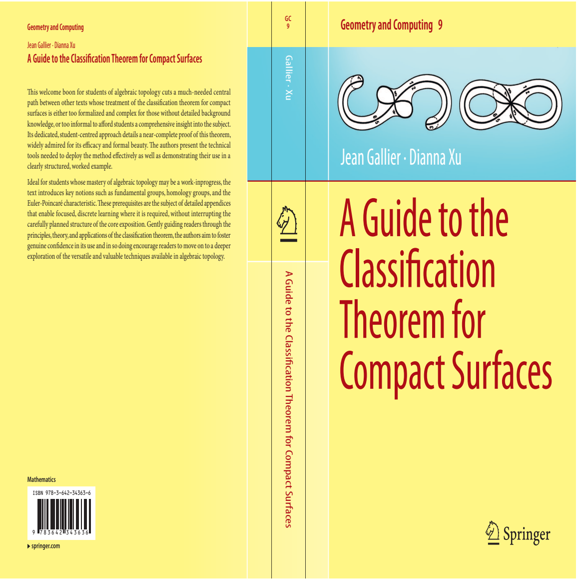

Geometry and Computing
Jean Gallier · Dianna Xu
A Guide to the Classification Theorem for Compact Surfaces
This welcome boon for students of algebraic topology cuts a much-needed central path between other texts whose treatment of the classification theorem for compact surfaces is either too formalized and complex for those without detailed background knowledge, or too informal to afford students a comprehensive insight into the subject. Its dedicated, student-centred approach details a near-complete proof of this theorem, widely admired for its efficacy and formal beauty. The authors present the technical tools needed to deploy the method effectively as well as demonstrating their use in a clearly structured, worked example.
Ideal for students whose mastery of algebraic topology may be a work-inprogress, the text introduces key notions such as fundamental groups, homology groups, and the Euler-Poincaré characteristic. These prerequisites are the subject of detailed appendices that enable focused, discrete learning where it is required, without interrupting the carefully planned structure of the core exposition. Gently guiding readers through the principles, theory, and applications of the classification theorem, the authors aim to foster genuine confidence in its use and in so doing encourage readers to move on to a deeper exploration of the versatile and valuable techniques available in algebraic topology.
Mathematics
ISBN 978-3-642-34363-6
springer.com
GC 9
Gallier · Xu
A Guide to the Classification Theorem for Compact Surfaces
Geometry and Computing 9
Jean Gallier · Dianna Xu
A Guide to the Classification Theorem for Compact Surfaces
Springer

Main steps of the proof

Sketch of the steps of the proof of Theorem 3. This proof follows Seifert and Threlfall (1934), which is basically Brahana's original proof (1921).

Main steps of the proof

Sketch of the steps of the proof of Theorem 3. This proof follows Seifert and Threlfall (1934), which is basically Brahana's original proof (1921).

Variants of this proof appear in Ahlfors and Sario, and Massey.

Main steps of the proof

Sketch of the steps of the proof of Theorem 3. This proof follows Seifert and Threlfall (1934), which is basically Brahana's original proof (1921).

Variants of this proof appear in Ahlfors and Sario, and Massey.

Step 1. Elimination of strings aa^{-1} in boundaries.

Main steps of the proof

Sketch of the steps of the proof of Theorem 3. This proof follows Seifert and Threlfall (1934), which is basically Brahana's original proof (1921).

Variants of this proof appear in Ahlfors and Sario, and Massey.

Step 1. Elimination of strings aa^{-1} in boundaries.

Step 2. Vertex Reduction.
The purpose of this step is to obtain a cell complex with a single vertex.

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

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

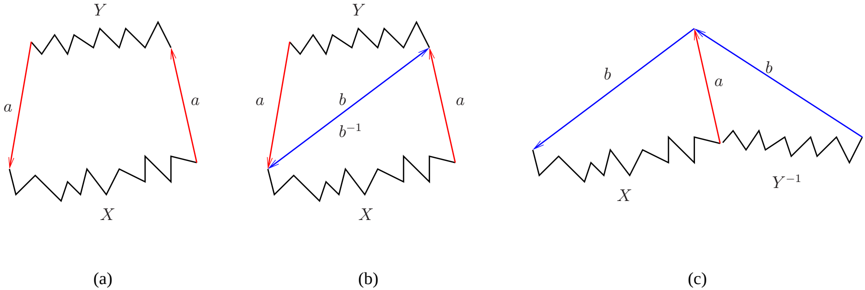

Figure : Grouping the cross-caps

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.

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.

Step 5. Transformation of handles into cross-caps.
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.

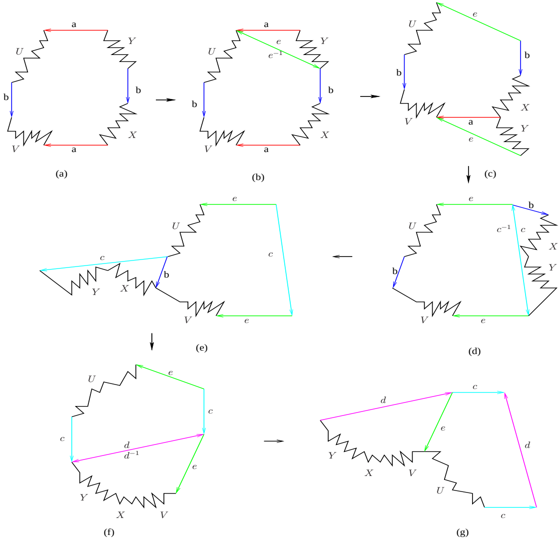

Figure : Grouping the handles

Fundamental Groups of Surfaces

For an orientable cell complex K (of type (I)), the fundamental group $\pi_1(K)$ is the group presented by the generators $\{a_1, b_1, \ldots, a_g, b_g\}$, and satisfying the single equation

$$a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} = 1.$$

When $g = 0$, it is the trivial group reduced to 1.

Fundamental Groups of Surfaces

For an orientable cell complex K (of type (I)), the fundamental group $\pi_1(K)$ is the group presented by the generators $\{a_1, b_1, \ldots, a_g, b_g\}$, and satisfying the single equation

$$a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} = 1.$$

When $g = 0$, it is the trivial group reduced to 1.

For a nonorientable cell complex K (of type (II)), the fundamental group $\pi_1(K)$ is the group presented by the generators $\{a_1, \ldots, a_g\}$, and satisfying the single equation

$$a_1 a_1 \cdots a_g a_g = 1.$$