THE CLASSIFICATION OF SURFACES

LENNART MEIER

> Everything should be made as simple as possible, but not simpler.
>
> Attributed to Albert Einstein

There are several textbook accounts on the classification of surfaces. I can recommend especially the ones in [Arm79], who presents a similar proof to ours, the informal presentation in [Zee66] and [Mas91], who presents essentially the most traditional proof. We present a different proof though, which is due to Thomassen [Tho92] and has not appeared in a textbook yet. We expand on some points that Thomassen just briefly sketches, in particular the well-definedness of attaching handles. The epigraph refers to the fact that the classification of surfaces contains some subleties, which are not equally well-treated in all sources. I hope that I managed to do better, though this document will certainly contain inaccuracies and gaps as well.

1. Surfaces

The aim of this series of lectures is to understand and classify surfaces in the following sense:

Definition 1.1. A *surface* is a connected compact 2-dimensional manifold (without boundary).

If we allow boundary, we speak explicitly of a surface with boundary, but retain that they are compact and connected. Note that the connectivity is only for convenience as every 2-dimensional manifold is the disjoint union of its connected components (exercise).

What are some examples of surfaces? Some are explicitly embedded into $\mathbb{R}^3$, like our good friend, the two-sphere. We can also hint at such an embedding by drawing a picture.[1]

 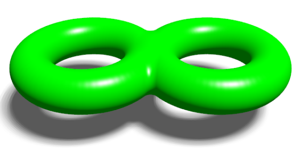 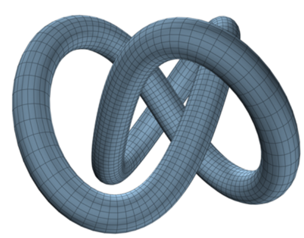

The number of "holes" in the surfaces is often called its *genus*. But in the moment it is not clear at all that this is a property of the manifold itself and not just of the embedding or the picture – and how many holes do you count in the rightmost picture anyhow?

[1]The rightmost of the pictures is taken from `https://forum.processing.org/two/discussion/23313/how-do-i-make-a-torus-knot`, the other two from wikicommons.

Another way to construct surfaces is to take a polygon and to glue pairs of side such that every side occurs in exactly one pair. It is not too hard to see that this gives a manifold (see e.g. [Lee11, Proposition 6.4]). An example is the following:

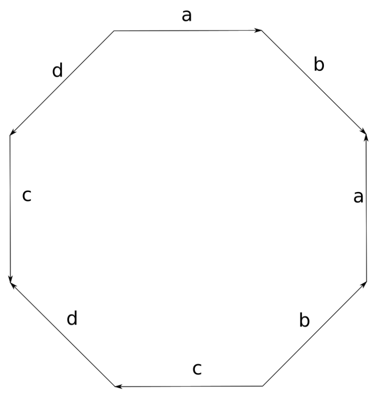

The letters indicate, which pairs of sides we glue. The arrows indicate, which homeomorphism of a side with the standard interval $[0, 1]$ we choose and hence by which homeomorphism we glue the two sides.

How to see whether this example is homeomorphic to one of the surfaces above? We can first test whether they have isomorphic fundamental groups. The surface obtained by glueing an octagon as above has fundamental group

$$\langle a, b, c, d \,|\, aba^{-1}b^{-1}cdc^{-1}d^{-1} \rangle.$$

The fundamental group of the genus two surface above can be computed via van Kampen. If we cut in a curve in the middle, one can see that the fundamental groups of the two pieces are both free on two generators (as they are tori without an open disk and thus homotopy equivalent to a wedge of two circles). We thus obtain a presentation with generators a, b, c, d and one relation: $aba^{-1}b^{-1} = c^{-1}d^{-1}cd$ (Note that our arguments in this computation are slightly vague as we defined the genus-2 surface above only by picture and not by a precise mathematical definition.)

The fundamental groups are thus isomorphic, but are the surfaces homotopy equivalent or even homeomorphic? The answer will be yes. We will indeed be able to completely classify surfaces up to homeomorphism. But first we will do as a warm-up the case of curves.

2. Warm-up: Classification of curves

As a warm-up we classify compact 1-dimensional manifolds (without boundary). We easily reduce to the connected case.

Definition 2.1. A *curve* is a 1-dimensional compact connected manifold without boundary.

Lemma 2.2. *Every curve C is homeomorphic to (the geometric realization of) a graph.*

Proof. Choose around every point of C a neighborhood that is homeomorphic to a closed interval. By compactness, finitely many suffices, and we call them A_i. Let the *complexity* of this cover be the number of pairs of intervals that overlap in their interior. We denote the complexity of C by n.

If the complexity is zero, then C is homeomorphic to a graph. We argue by induction and assume that every curve C with a cover of complexity smaller than n is homeomorphic to a graph. Consider two overlapping intervals A_i and A_j. If $A_i \subset A_j$, delete A_i. Else, replace A_j by $\overline{A_j \setminus (A_i \cap A_j)}$, which is still an interval. In either case, the new cover has lower complexity. □

Remark 2.3. For a version of this result without compactness assumption see [Lee11, Theorem 5.10].

Proposition 2.4. *Every curve C is homeomorphic to S^1.*

Proof. By the previous lemma, there is a graph structure on C, which we will fix. As C is a 1-dimensional manifold, removing v from a small neighborhood of a vertex v results in a space with two connected components. Thus, every vertex must have valence 2, i.e. must be contained in exactly two edges (which might be identical if we allow loops).

We claim that C is as a graph isomorphic to a standard n-gon D_n. As (the geometric realization of) D_n is homeomorphic to S^1, this implies the result.

We argue by induction on the number n of edges, the case of exactly one edge being clear as the only graph with just one edge has one vertext and one loop is thus isomorphic to D_1. If we have more than one edge, pick an edge e and consider the quotient C/e. This is still a curve and thus induction shows that $C/e \cong D_{n-1}$. If we cut D_{n-1} open along the vertex correspond to e and insert an edge, we obtain both C and D_n and thus these have to be isomorphic. □

This foreshadows our procedure in the significantly more complicated case of surfaces:

(1) Find a combinatorial structure on our object.
(2) Simplify the combinatorial structure by an inductive procedure and obtain the original surface by performing these simplifications backwards.

3. Triangulations

The precise combinatorial structure for surfaces we need to consider are *triangulations.* Combinatorially, triangulations are given by simplical complexes.

Definition 3.1. A *triangulation* of a surface S is a homeomorphism $S \to K$ to a space K that we obtain as follows: Consider a disjoint union of triangles (i.e. Δ^2) and a set of pairs of directed edges, which we glue together. We call the images of vertices, edges and triangles by the same name in K. We demand that

- every edges in K has two distinct endpoints and for every two vertices there is at most one edge between them, and
- every triangle has three distinct sides and for every three vertices there is at most one triangle spanned by them.

Such a space K is called a (2-dimensional) *simplicial complex.* If we leave out the two last conditions (i.e. K is just glued from triangles along pairs of directed edges) we call K a Δ-*complex.*[2]

Combinatorially, a 2-dimensional simplicial complex can be encoded in a set of vertices V, a set E (of edges) of two-elements subsets of V and a set T (of triangles) of three-elements subsets of V such that every two-element subset of some triangle $t \in T$ is in E. This combinatorial datum is sometimes called an *abstract simplicial complex.* One can

[2]Simplicial complexes and Δ-complexes are actually more general concepts that also exist in arbitrary dimensions. (We also demanded that every vertex and every edge is part of a triangle and that every edge is contained in at most two triangles, which is usually not part of the definition of a simplicial or Δ-complex.) Except in a short outlook, we will only consider the kind of simplicial complexes and Δ-complexes as above.

similarly define higher dimensional version, where we glue a disjoint set of higher dimensional simplices Δ^n together along faces. Here, Δ^n stands for the convex hull of the unit vectors in $\mathbb{R}^{n+1}$, i.e. an n-dimensional analogue of triangle and tetrahedron. Abstractly, a higher-dimensional simplicial complex is encoded in subset $K_i \subset V^i$. (See [Lee11] for more details.)

Lemma 3.2. *Every edge in a triangulation of a surface is contained in two triangles (in other words: every side of a triangle is glued to another side). Moreover, every vertex has valence at least* 3*, where valence refers to the number of edges adjacent to the vertex.*

Proof. Assume that there is an edge e that is only part of one triangle and let x be a point in the interior of e. Then x has arbitrary small neighborhoods V in S such that $V \setminus \{x\}$ is simply-connected, even contractible. In contrast, let U be a neighboorhod of x that is homeomorphic to $\mathbb{R}^2$ (which exists as S is a surface). Take a neighborhood V such that $V \setminus \{x\}$ is simply-connected and $V \subset U$. Then take a small loop around x that is non-nullhomotopic in $U \setminus \{x\}$ and that is contained in V. But as it is nullhomotopic in $V \setminus \{x\}$, we arrive at a contradiction. Thus every edge must be part of two triangles.

Clearly, every vertex must have valence at least 2. Assume that v is a vertex of valence 2 and let w and u be the two other endpoints of the adjacent edges. Let e be one of the edges adjacent to v and let Δ_1 and Δ_2 be the two triangles adjacent to e. Each of these triangles has two edges adjacent to v and we see that both of them in Δ_1 are glued to the corresponding ones in Δ_2 as v has valence 2. Thus, Δ_1 and Δ_2 are two triangles with the vertices u, v and w, in contradiction with the definition of a triangulation. □

The tetrahedron and the octahedron are examples of simplicial complexes that are homeomorphic to the sphere S^2. Indeed: Embed the tetahedron T into $\mathbb{R}^3$ such that 0 is in its interior and T lies inside S^2. The projection $\mathbb{R}^3 \setminus \{0\} \to S^2$ restricts to a continuous bijection $T \to S^2$ that is automatically a homeomorphism as T is compact and S^2 is Hausdorff. The argument for the octahedron is analogous.

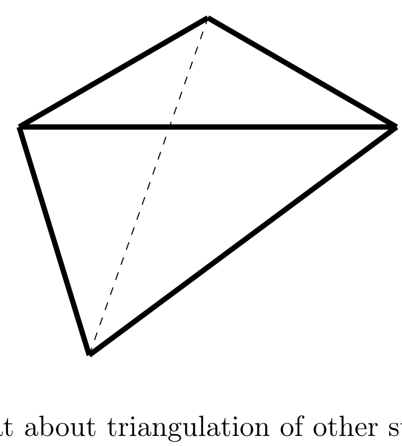

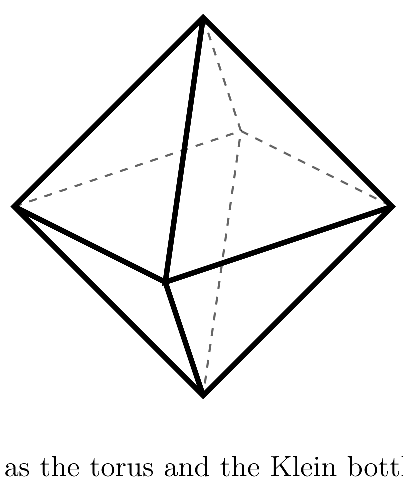

What about triangulation of other surfaces such as the torus and the Klein bottle? One could think in error that these are examples of triangulation of these surfaces.

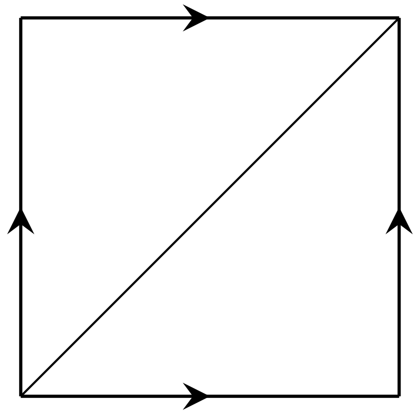

The torus

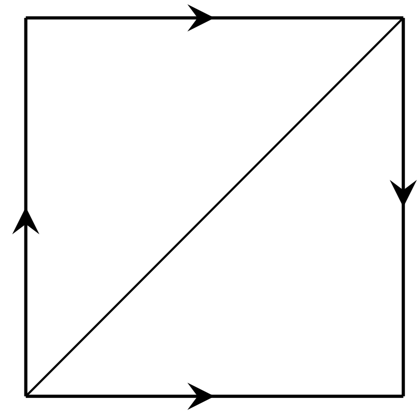

The Klein bottle

But note that this violates the condition that every edge must have two distinct endpoints. For example in the case of the torus, all four corners are identified to one point after glueing. Thus we only have a Δ-complex and not a simplicial complex. To obtain a simplicial complex (and hence a triangulation) one has to subdivide the triangles to obtain a triangulation. A systematic procedure is the *barycentric subdivision*, which has additional vertices the midpoints of all sides and the barycenter of the triangle.

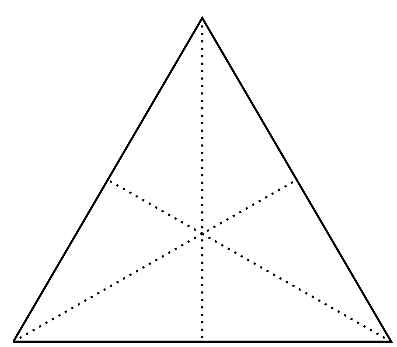

If we choose the barycentric subdivision of the two triangles in the torus, we still do not obtain a triangulation (exercise: which edges have the same endpoints?). But if we subdivide all the triangles again barycentrically, it provides a triangulation. Actually this works for every space glued from triangles along sides that if we doubly barycentrically subdivide all its triangles, it yields a triangulation. In the case of the torus, this triangulation has 72 triangles. In contrast, there are simpler triangulation with only 14 triangles, of which we depict an example.

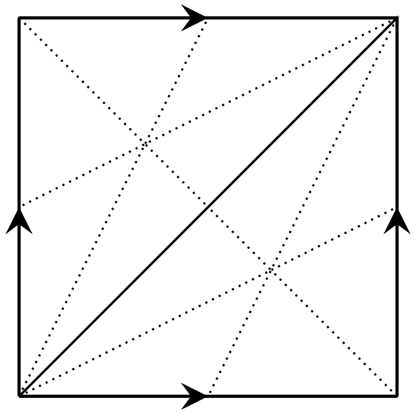

The barycentric subdivision

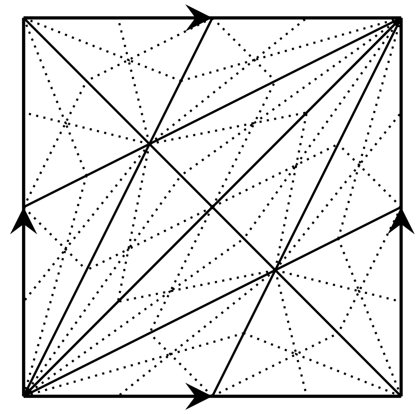

The double-barycentric subdivision

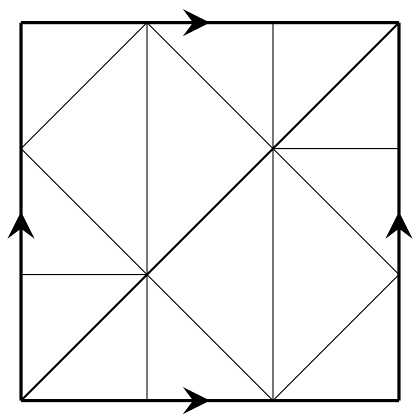

A minimal triangulation of the torus

While we have proven above the easy statement that every 1-dimensional manifold has a "triangulation", the following 2-dimensional analogue is harder and we will use it without proof (see [Tho92] for a modern, self-contained proof).

Theorem 3.3 (Radó, 1925)**.** *Every surface has a triangulation.*

Moise (1952) shows that every compact 3-dimensional manifold has a triangulation as well (i.e. it admits a homeomorphism to a 3-dimensional simplicial complex). It might come as a surprise that for every $n \geq 4$, there is a compact n-dimensional manifold (without boundary) that does not have a triangulation. The latter results are very deep. For dimension 4 it follows from results of Freedman (1982) and Casson (1990) – this is actually part of the work Freedman won the fields medal for. For dimensions at least 5 the result is due to Manolescu (2013). In contrast, it was proved much earlier by Cairns (1935) and Whitehead (1940) that every *differentiable* manifold has a triangulation. Thus, these counterexamples in dimensions at least 4 are truly weird: topological manifolds without differentiable structure or triangulation. Luckily, in dimension 2 we do not have to face such problems.

An *isomorphism* between two triangulations of two surfaces consists of a bijection $f\colon V \xrightarrow{\cong} V'$ of the sets of vertices such that we have for the sets of edges and triangles the equalities $f(E) = E'$ and $f(T) = T'$. Given two triangulated manifolds $M \cong K$ and $N \cong K'$ and an isomorphism $K \cong K'$, we obtain a composed *combinatorial isomorphism* $M \cong K \cong K' \cong N$. But there are other "combinatorial" ways M and N can be homeomorphic. We may namely subdivide a triangle into sub-triangles. We allow any subtriangulation, where the new triangles are (affine) linearly imbedded into the standard 2-dimension simplex. Besides the example already depicted above, the most important example for us the following simple subdivision, also called the *stellar subdivision*:

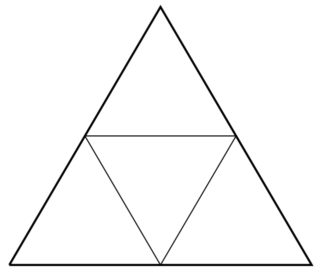

Definition 3.4. A subdivision of a simplicial complex or more generally a Δ-complex consists of subdivision of all constituting triangles such that the subdivisions match on the glued edges. Two Δ-complexes are *combinatorially homeomorphic* if they have a common subdivision, i.e. if there exist subdivisions of the two complexes that are combinatorially isomorphic. Two triangulated surfaces are combinatorially homeomorphic if the associated simplicial complexes are combinatorially homeomorphic.

It is not hard to see that combinatorial homeomorphisms are indeed homeomorphisms.

Lemma 3.5. *Being combinatorially homeomorphic is an equivalence relation on the class of Δ-complexes (and hence also on the subclass of simplicial complexes).*

Proof. Symmetry and reflexivity are clear. Now suppose K and L have a common subdivision K' and L and M have a common subdivision L'. We need to show that K' and L' have a common subdivision again. As they are both (up to isomorphism) subdivision of L, it suffices to see that two subdivisions of the same Δ-complex have a common subdivision. This follows from the fact that the intersection of two linearly embedded triangles into a triangle has a triangulation (as illustrated in the picture below).

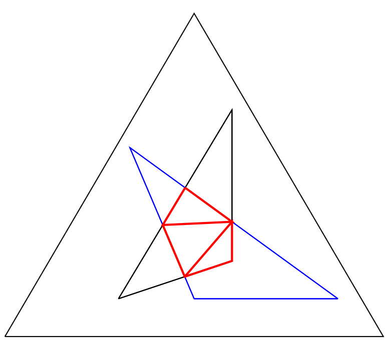

□

One special case of a combinatorial homeomorphism is the one between the two obvious triangulations of a quadrilateral.

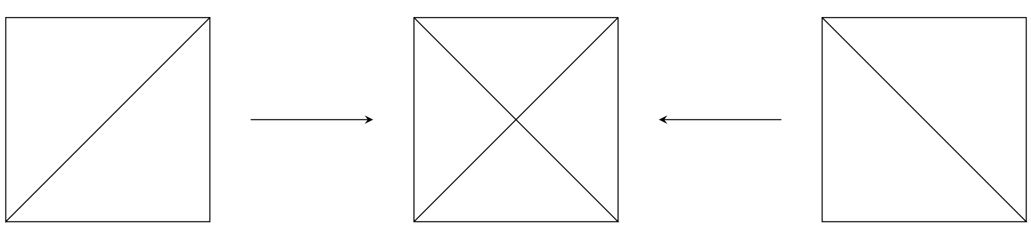

In practice a combinatorial homeomorphism will for us be a composition of "moves", where we subdivide (or take the inverse of a subdivision) or change the triangulation of a quadrilateral.

For more information about triangulations and simplicial complexes we refer to the books by Armstrong [Arm79] and [Lee11]. For Δ-complexes see also [Hat02].

4. Statement of the classification and orientability

The classification will be of the form that every surface is homeomorphic to exactly one in a list of standard surfaces. First we have the list of surfaces with "g holes". For $g = 0$, this is just S^2 (with, say, the triangulation given by the tetrahedron), for $g \geq 1$ we glue a $4g$-gon in the following scheme to obtain a surface S_g:

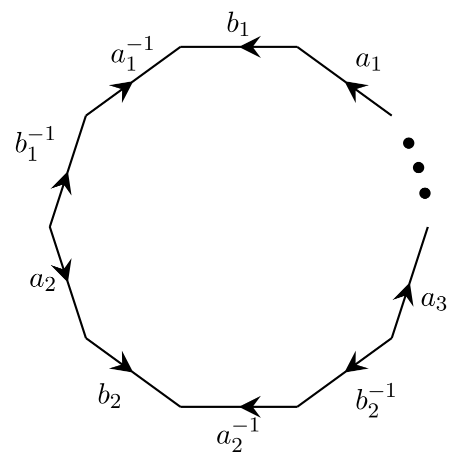

But there are also other surfaces, not fitting in this scheme. For example, $\mathbb{RP}^2$ and the Klein bottle are not homeomorphic to any of these, as we will see. One difference is that the surfaces S_g are *orientable* while the the surface $\mathbb{RP}^2$ and the Klein bottle are *not orientable*. Intuitively, orientable means that we can give a disk in our surface an orientation (i.e. a notion what it means to rotate clockwise or counterclockwise) and if we move our disk along a path to the same disk again, this notion cannot change. If we can embed a Möbius strip into our surface, our surface is thus not orientable. Combinatorially, an orientation means that every triangle in our triangulation gets a sense of clockwise orientation such that the corresponding orientations of the edges are opposite for two adjacent triangles. See the discussion in [Arm79, Section 7.2] for details.

There are more non-orientable surfaces than just $\mathbb{RP}^2$ and the Klein bottle, namely the surfaces N_g, which we obtain by glueing a $2g$-gon in the following scheme:

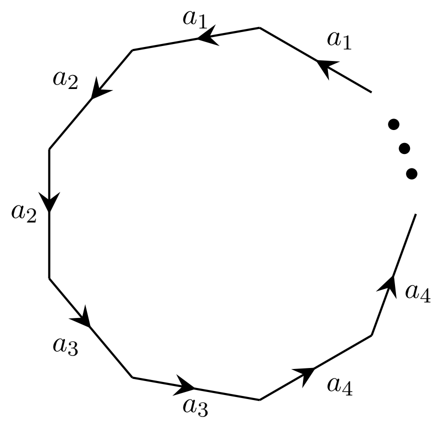

As an exercise find the embedded Möbius strip inside these surfaces.

The surfaces N_1 and N_2 are $\mathbb{RP}^2$ and the Klein bottle. We call the collection of S^2, S_g and N_g (with $g \geq 1$) the *standard surfaces.* We sometimes set $S_0 = S^2$.

Theorem 4.1. *Every surface is homeomorphic to exactly one of the standard surfaces.*

What we will actually prove is a slightly different theorem, but it will directly imply Theorem 4.1 together with Theorem 3.3.

Theorem 4.2. *Every triangulated surface is combinatorially homeomorphic to one of the standard surfaces.*

Remark 4.3. The classification of (oriented) surfaces was first stated by Möbius in 1861, but his proof precedes the modern notions of topological spaces and manifolds by several decades and is thus not rigorous by modern standards. The first rigorous proof of the classification of triangulated surfaces was maybe given by Dehn–Heegard (1907). Together with Radó's theorem, one can say that the history of the classification of surfaces, from its first statements to a complete proof, stretches from 1861 to 1925.

One thing we can directly prove is that none of the standard surfaces are homeomorphic (or even homotopy equivalent) to each other and so that every surface can be homeomorphic to at most one of these. We use the fundamental group. We have

$$\pi_1(S_g) \cong \langle a_1, b_1, \ldots, a_g, b_g \,|\, (a_1 b_1 a_1^{-1} b_1^{-1}) \cdots (a_g b_g a_g^{-1} b_g^{-1}) \rangle.$$

In contrast, we have:

$$\pi_1(N_g) \cong \langle c_1, \ldots, c_g \,|\, c_1^2 \cdots c_g^2 \rangle.$$

Are any of these groups isomorphic? They seem different, but proving whether two non-abelian groups are isomorphic is a bit tricky. (In general, there is not even an algorithm to decide whether a given group presentation defines the trivial group!) But for abelian groups, it is usually much easier. Thus, we introduce the *abelianization* G^{ab} of a group G. It comes with a group homomorphism $\phi\colon G \to G^{ab}$. Formally speaking, it is characterized by the following universal property: Every group homomorphism $G \to H$ into an abelian group factors uniquely as a composition $G \to G^{ab} \to H$. Alternatively it can be characterized as the maximal abelian quotient. Given a presentation, we can construct it by just adding the relations that all generators commute. E.g. in case of $\pi_1(N_g)$ we would add the relations $c_i c_j = c_j c_i$. Note that if two groups are isomorphic, also their abelianizations are isomorphic (as follows from the universal property).

We directly see that in the abelianization of $\pi_1(S_g)$ the relation becomes trivial and thus $\pi_1(S_g)^{ab}$ is just the free abelian group on the generators $a_1, b_1, \ldots, a_g, b_g$, i.e. $\mathbb{Z}^{2g}$. The abelianization of $\pi_1(N_g)$ is isomorphic to the free abelian group of the generators $c_1, \ldots, c_g$ modulo the relation $2c_1 + \cdots + 2c_g$ (we switched to addtive notation as we are in an abelian group now). We can change the basis of the free abelian group to $e_1, \ldots, e_g$ with $e_i = c_i$ for $1 \leq i \leq g-1$ and $e_g = c_1 + \cdots + c_g$, where the relation reads $2e_g = 0$. Thus, $\pi_1(N_g)^{ab} \cong \mathbb{Z}^{g-1} \oplus \mathbb{Z}/2$. We see that no two of the groups $\pi_1(S_g)$ and $\pi_1(N_g)$ are isomorphic and thus no two of the S_g and N_g are homotopy equivalent or even homeomorphic. (See also [Lee11, p. 264-267].)

5. Handles and cross-caps

There is also a different way of building surfaces than via polygonal gluings. To gain some intuition, we will first describe this construction without triangulations and only introduce triangulations on them only later.

We can obtain a surface of genus g by attaching g handles to a sphere. Here, *attaching a handle* means the following: Embed two disks disjointly via maps $\phi_1, \phi_2 \colon D^2 \to S$ into a surface S – if S is orientable, we demand that the disks have opposite orientation.[3] Remove the interiors of the two disks to obtain a space S' and attach a cylinder at the resulting two boundary circles. More precisely, we obtain this as the quotient space of $S' \coprod S^1 \times I$, where we identify $(x, 0) \in S^1 \times I$ with $\phi_1(x) \in S'$ and $(x, 1)$ with $\phi_2(x)$. A *cross-handle* is the same, but the orientations of the disks agree. If we have a surface embedded into $\mathbb{R}^3$ and attach a handle, it can still be embedded into $\mathbb{R}^3$, while a cross-handle cannot be embedded anymore, but would need some self-intersection.[4] It is not hard to see that the resulting space of attaching a (cross-)handle is still a surface. For pictures of this process we refer to the book [Arm79] and the notes [Zee66].

Remark 5.1. Let ϕ_1 and ϕ_2 be as above and D^2_r be a disk of radius r. Then an alternative way to obtain the surface that results from attaching a (cross-)handle along the ϕ_i is the following: Remove the interiors of the disks $\phi_i(D^2_{\frac{1}{2}})$ from S to obtain a new space S''. Now identify $\phi_1(x)$ with $\phi_2(x)$ for $x \in \partial D^2_{\frac{1}{2}}$.

Here we use that we obtain the cylinder by glueing $D^2 \setminus \mathring{D}^2_{\frac{1}{2}}$ with itself along the inner boundary circle.

Attaching a *cross-cap* means to remove the interior of an embedded disk and identifying opposite points on the boundary. It is not hard to see that the resulting space of attaching a cross-cap is still a surface. Again we refer for pictures to [Arm79] and [Zee66].

Next we want to explain versions of this for a triangulated surfaces, where we triangulate the resulting new surface again.

In this setting, instead of choosing embedded disks to attach a (cross-)handle at, we will choose two triangles in our triangulation that are not connected by any edge. We remove the interiors of the triangles and identify their corresponding sides. This produces

[3] We have not discussed orientations of non-triangulated surfaces and we will not actually use them either. But one way to define this is to say that that the two embeddings $\phi_1, \phi_2 \colon D^2 \to S$ are not isotopic, i.e. there is no homotopy (H_t) from ϕ_1 to ϕ_2 such that H_t is an embedding for each $t \in I$ – such a homotopy is called an *isotopy*.

[4] This is a fact we actually won't prove here. One approach to see this is that after attaching a cross-handle the surface becomes non-orientable and non-orientable (closed) surfaces cannot be embedded into $\mathbb{R}^3$.

a triangulation on the new surfaces with the three less vertices, two less triangles and three less edges. (See the remark above why we call this attaching a (cross-)handle.)

For attaching cross-caps, we pick a vertex v of valence $2g$ ($g \geq 3$)[5] and consider the disk that is the $2g$-gon obtained from all the triangles touching v. Removing the interior and identifying opposite sides produces a cross-cap.

6. Surgeries and Euler characteristic

Attaching a handle or a cross-cap makes our surface topologically more complicated (if we say that the sphere is the simplest surface). We can also go in the opposite direction to make our surface simpler.

Suppose there is a path $\gamma = v_1v_2v_3v_1$ on a triangulated surface S (along edges) that does not separate S, i.e. its complement is still connected. There are now two cases, depending on whether γ is *orientation-preserving* or *orientation-reversing.* In the former case, a neighborhood of the curve will look like a cylinder, while in the latter case it looks like a Möbius strip. (I encourage you to try to cut a Möbius strip along the middle circle to see that this cut does not separate it into two components.)

More precisely, we distinguish these two cases by considering what happens if we *cut* the surface S along γ. This process can be described as follows: S is glued from a number of triangles along their boundaries. We look at the triangulated surface S' (with boundary!) where we leave out the gluings along the edges of γ. The resulting boundary is a graph with six edges and two properties: Every two vertices are connected by at most one edge and every vertex has valence two. This shows that the boundary consists either of two triangles or of one hexagon. In the former case we say that γ is orientation-preserving, in the latter that γ is orientation-reversing.

If S' has two boundary triangles, we produce a surfaces S'' by gluing two triangles along these boundary curves. In the other case, we glue a hexagon (as consisting of six triangles) into the boundary curve. The resulting surface S'' does not have a boundary anymore is said to be obtained via a *surgery along* γ.

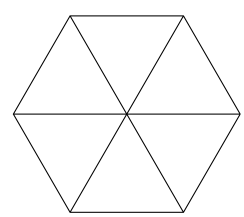

A hexagon

It is direct from the definition that if γ was orientation-preserving, we can get S back from S'' by attaching a (cross-)handle and if it was orientaton-reversing by attaching a cross-cap.

Now we need a measure to say that the surface became simpler. This is the Euler characteristic. If the triangulation T of S has n vertices, e edges and t triangles, its Euler characteristic is defined to be $t - e + n$. We denote it either by $\chi(T)$ or by $\chi(S)$ when the triangulation is understood. (We will see later that the Euler characteristic only depends on the homeomorphisms or even homotopy type of S, not on the specific triangulation. But in the moment, we do not need to use this fact yet.)

If we do a surgery along a orientation-preserving path, we gain three vertices, three edges and two triangles. Thus the Euler characteristic is raised by 2.

If we do a surgery along a orientation-reversing path, we gain four vertices, nine edges and six triangles. Thus the Euler characteristic is raised by 1.

Lemma 6.1. *If S is a triangulated surface, then $\chi(S) \leq 2$.*

[5]This means that there are $2g$ edges adjacent with v.

Proof. Look at the graph G of vertices and edges and choose a maximal tree T inside of it. Now consider the *dual graph* D of T, i.e. its vertices are the triangles of S and its edges are those edges not in T.

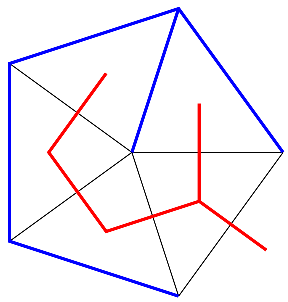

The maximal tree T is marked in blue

The dual graph D is marked in red

We claim that the dual graph D is connected and actually more generally the dual graph D' of every tree T' inside G is connected. We prove this by induction on the number of vertices n in T'. Assume we have proven it for all trees with less than n vertices and choose a vertex v in T' just adjacent to one edge e. We know by induction that the dual graph of $T' \setminus \{e\}$ is connected. The only way D' can not be connected is if the two triangles Δ_1 and Δ_2 adjacent to e are not connected in D'. But none of the other edges $e_1, \ldots, e_k$ (ordered clockwise) adjacent to v are in T; thus $e_1, \ldots, e_k$ forms a path in D' connecting Δ_1 and Δ_2. Thus the dual graph is indeed connected.

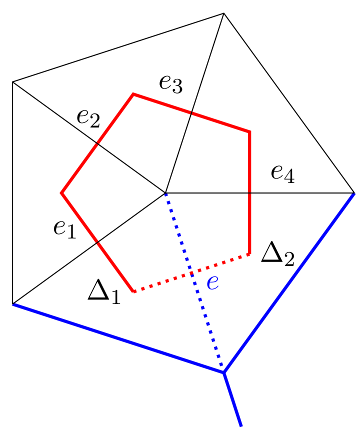

An inductive argument shows that every tree and in particular T has Euler characteristic 1, while every connected graph and thus D has Euler characteristic at most 1. As T contains all vertices, $\chi(T) + \chi(D) = \chi(S)$ and we see $\chi(S) \leq 1 + 1 = 2$. □

In particular, we can only do finitely many surgeries at one surface and none at one of Euler characteristic 2.

Remark 6.2. It is an important observation that the Euler characteristic does not change by subdivision. This is very easy to see for the only two types of subdivision we use, namely the stellar subdivision (where we add one vertex, three edges and two triangles) and the one moving from one triangulation of the quadrilateral to another.

In general, there is an induction argument on the number of edges, which is easiest to set up not just considering triangulations, but arbitrary graphs inside the triangle (with the count of triangles replaces by the count of faces, i.e. connected regions in the complement of the edges and vertices). If you want to see the details of the proof, google *Euler's formula for graphs*.

7. The main proof

In this section, we will prove a weak version of our main classification theorem, but which will be sufficient with some extra work to obtain Theorem 4.2.

Theorem 7.1. *Every triangulated surface S is combinatorially homeomorphic to a surface that is obtained from a tetrahedron by iteratively attaching (cross-)handles and cross-caps (and possibly subdividing in this process). The Euler characteristic of S equals $2-2h-c$, where h is the number of (cross-)handles and c the number of cross-caps attached.*

Proof. Start with any triangulation T of our surface S. Let $\chi(T)$ be the Euler characteristic of this triangulation of S. In case that (S,T) does not satisfy the conclusion of the theorem, we call it a *counterexample*. We assume that S is a counterexample with the following properties:

(1) The Euler characteristic $\chi(T)$ is maximal among all counterexamples (we can assume this as $\chi(T) \leq 2$ by Lemma 6.1),
(2) The number of vertices in T is minimal among all counterexamples subject to (1),
(3) The minimal valence of a vertex of T is minimal among all counterexamples subject to (1) and (2), where valence still refers to the number of edges adjacent to the given vertex.

Let v be a vertex in T with minimal valence. By Lemma 3.2, the valence must be at least 3. If the valence is 3, we can remove v and the three edges adjacent to v as in the picture and obtain a triangulation T' such that (S,T) is combinatorially homeomorphic to (S,T'). (Here we are using that T is not already a tetrahedron.)

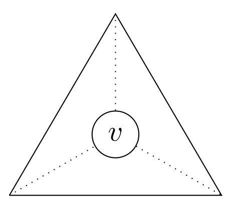

Deleting a vertex of valence 3

Moreover, we have $\chi(T) = \chi(T')$, but the latter triangulation has less vertices than T. This is in contradiction with (2).

Thus, the valence of v is at least 4. Call the adjacent vertices $v_1, \ldots, v_q$. Assume first that there is an i such that v_i and v_{i+2} (cyclically counted so that e.g. $v_{q+2} = v_2$) are not connected by an edge. The following picture demonstrates how to change the triangulation to a combinatorially homeomorphic one, where the vertex corresponding to v has lower valence, in contradiction to (3).

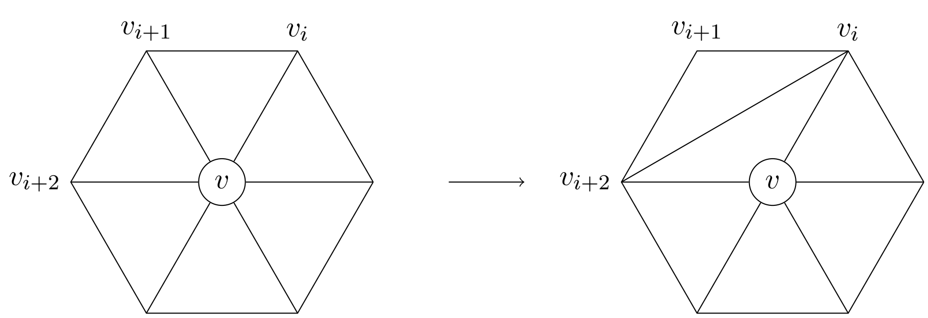

Thus, we can assume that v_i and v_{i+2} are connected by edges for all $i = 1, \dots, q$. We claim that the path vv_1v_3 does not separate the surface. Indeed, the intersection of a small neighboorhood of it with $S \setminus vv_1v_3$ has at most two components and because of the cyclic ordering v_2 and v_4 must lie in different ones of them if there are two. Every path on S crossing vv_1v_3 must cross this neighborhood and it is thus enough to observe that v_2 and v_4 are connected by an edge as these two points are connected to every other point in the neighborhood.

Now we do a surgery along the path vv_1v_3. Our discussion about surgeries implies that after the surgery, the Euler characteristic of the triangulated surface becomes larger and thus the resulting surface S' can be obtained via attaching (cross-)handles and crosscaps from the sphere (with its tetrahedron triangulation) by induction. Moreover, we have seen that S can be obtained from S' via attaching a (cross-)handle or cross-cap. Thus S cannot be a counterexample. □

This already implies that every surface of Euler characteristic 2 is combinatorially homeomorphic to the tetrahedron. We call such surfaces just *spheres*.

8. Well-definedness of attaching (cross-)handles and cross-caps

We prove the classification of surfaces with boundary only in the simplest cases. First, we will need a lemma.

Lemma 8.1. *Let S be a connected compact surface with boundary. Then $\chi(S) \leq 1$.*

Proof. Let $c_1, \dots, c_k$ be the boundary cycles of S such that c_i has n_i edges each. Let D_n be an n-gon that is triangulated with one inner vertex and n triangles. Thus $\chi(D_n) = 1$. We can glue S and $D_{n_1}, \dots, D_{n_k}$ by glueing c_i with the boundary of D_{n_i}. The resulting surface does not have boundary anymore and will have Euler characteristic $\chi(S) + k \leq 2$ (by Lemma 6.1). Thus $\chi(S) \leq 1$. □

Proposition 8.2 (Disk and annulus theorem)**.** *Let S be a triangulated surface with boundary.*

If S has one boundary component and Euler characteristic 1*, the surface is combinatorially homeomorphic to a disk (more precisely, to a triangle). We call such a surface a* combinatorial disk*. If S has two boundary components and Euler characteristic* 0*, the surface is combinatorially homeomorphic to an annulus (as in the picture below). We call such a surface a* combinatorial annulus.[6]

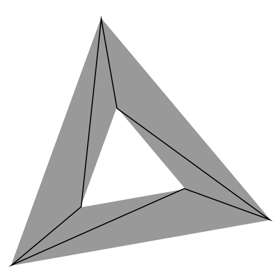

A triangulated annulus

[6]The proof actually gives the following more precise statement: Given two combinatorial disks any combinatorial homeomorphism of their boundaries extends to a combinatorial homeomorphism of the combinatorial disks. For combinatorial annuli, any combinatorial homeomorphism of one of the boundary cycles of the first annulus to one of the boundary cycles of the second annulus extends to a combinatorial homeomorphism of the annuli. (But we are not allowed to prescribe combinatorial homeomorphisms of both cycles.) This will be the result we will actually use.

Proof. The proof is an adaption of the argument in Theorem 7.1. Let S be a surface of Euler characteristic 1 with one boundary cycle with n edges. Assume first that S has an inner vertex. Then there exists an inner vertex v of minimal valency. If the valency is 3, then we can remove v to obtain a triangulation with less inner vertices as in the proof above. If the valency is at least 4, we can consider all adjacent vertices $v_1, \dots, v_q$ (cyclically ordered). If not every v_i and v_{i+2} are adjacent, we can make a move to reduce the valency of v as in the proof above. We claim that not all v_i and v_{i+2} can be adjacent. Indeed, if both v_1 and v_3 are inner, we can make a surgery along the triangle vv_1v_3 as before and get a larger Euler characteristic, which is impossible (as the resulting surface will still have non-empty boundary). If v_1 or v_3 is not inner, we can attach to the boundary of S a "ring" so that all vertices on the boundary become inner; for a concrete ring, it is easy to see that the Euler characteristic of S does not change and so the same argument applies.

All in all we obtain by induction a triangulation without inner vertices of an n-gon (as there must be exaxctly one boundary curve). Choosing any edge that separates the n-gon lets us show by induction that all of these triangulations are combinatorially homeomorphic. Thus, all surfaces of Euler characteristic 1 with one boundary cycle with n edges are combinatorially homeomorphic. Subdivision shows that the triangle is combinatorially homeomorphic to such a surface for every n.

Now let S be a surface of Euler characteristic 0 with two boundary cycles of lengths m and n, respectively. Choose a path P of length l from one boundary cycle to the other. We can cut at this path to obtain new surface S' with just one boundary cycle and Euler characteristic 1, whose boundary is divided into segments of lenght m, l, n and l. The arguments above show that without any subdivision of the boundary, this is combinatorially homeomorphic to any other combinatorial disk with an analogous division of the boundary. Glueing along the paths of length l again, we obtain a surface that is combinatorially homeomorphic to S. Clearly, $l \geq 1$, and also $m, n \geq 3$ as else the resulting "triangulation" of S is not really a triangulation as there will be either a loop or two edges with the same endpoints. The following picture indicates part of a triangulation of a a combinatorial disk with the required segmentation of the boundary.

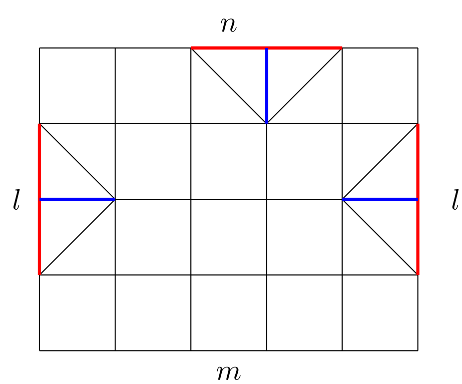

The (part of the) triangulation displayed is a subdivision of the same triangulation with the blue edges removed. By this removal, the red-marked pairs of edges are becoming one edge each. Thus, we have shown how to find a combinatorially homeomorphic triangulation with smaller l, m, n if $l > 1$ or m or n is bigger than 3. This way we can assume inductively up to combinatorial homeomorphism that $l = 1$ and $m, n = 3$ and as we just said, all triangulation with this property are combinatorially homeomorphic. □

Proposition 8.3. *Let S be a triangulated surface, possibly with boundary. Let D_1, D_2 be two simplicial subcomplexes that are combinatorial disks and do not touch the boundary. Then there is a combinatorial homeomorphism $S \to S$ restricting to a combinatorial homeomorphism $D_1 \to D_2$.*

Proof. Let D be a combinatorial disk inside S. Draw a "ring" R around D as in the picture.

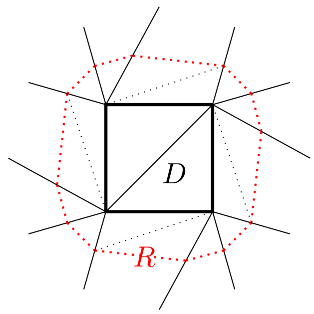

More precisely, we choose a point on every edge adjacent but not contained in D. As the vertices of D have a cyclic ordering and the edges adjacent to a vertex have one as well, these new points have a cyclic ordering as well. With respect to this cyclic ordering we connected every of these points to the next one, forming a polygon, which is the ring R. The resulting figure will not only contain triangles, but also quadrilaterals. Drawing diagonals in the latter produces a subdivision of the original triangulation.

The area A_D between R and D is a combinatorial annulus as it has two boundary components (the boundary of D and R) and the Euler characteristic is easily calculated to be zero. Let T be an arbitrary triangle in D. We claim that the area A_T between R and T is also a combinatorial annulus. Indeed: The area F enclosed by R has Euler characteristic $\chi(D) + \chi(A_D) = 1$ as the intersection of D and A_D has Euler characteristic zero. As A_T is F with one missing triangle, we see that $\chi(A_T) = 0$.

Thus A_D and A_T are combinatorially homeomorphic by the last proposition and this combinatorial homeomorphism φ can be chosen to be the identity on R. Moreover, D and T are combinatorially homeomophic by the last proposition as well and we can choose the combinatorial homeomorphism to be the restriction of φ on ∂D. Thus, these combinatorial homeomorphism glue to a combinatorial homeomorphism of S to S that is the identity outside of F and takes D to T.

We see that it suffices to find for two triangles combinatorial homeomorphisms that take one to the other. As two adjacent triangles T_1 and T_2 form together a combinatorial disk D, this is clear for adjacent triangles (using the composite of the combinatorial homeomorphism transporting T_1 to D and the one transporting D to T_2). It is easy to see that any two triangles are connected by a chain of adjacent triangles by just following a path between the two triangles. □

Corollary 8.4. *Attaching a (cross-)handle along two chosen (non-touching) triangles does not depend on the chosen triangles, only possibly on their orientations, up to combinatorial homeomorphism. Attaching a cross-cap does not depend on the chosen $2g$-gon up to combinatorial homeomorphism as well.*

Proof. For the first case, let T_1 and T_2 respectively T_1' and T_2' such triangles. Use the preceding proposition to produce a combinatorial homeomorphism that sends T_1 to T_1'. Thus we can assume $T_1 = T_1'$. Remove the interior of T_1 to obtain a surface with boundary. Then there is still a combinatorial homeomorphism sending T_2 to T_2', fixing the boundary (possibly not pointwise). Glueing T_1 in again shows in total that there is a combinatorial homeomorphism of our surface to itself sending T_1 to T_1' and T_2 to T_2'. This easily implies the first part.

The second part is similar. □

Remark 8.5. The arguments of this section also imply that we have more freedom in attaching (cross-)handles and crosscaps than previously stated. For a cross-cap we can take an arbitrary combinatorial disk with an even number of boundary edges, take out the interior and identify opposite sides. The triangulation in the interior does not matter. Similarly for (cross-)handles, where have only to ensure that both combinatorial disks have the same number of boundary edges. It might be necessary to subdivide afterwards to obtain a triangulation again.

Example 8.6. Choose the 4-gon on a tetrahedron consisting of two of the triangles. Attaching a cross-cap at this produce the usual rectangle with gluings that defines $\mathbb{RP}^2$. (We need to subdivide to provide a triangulation again)

Example 8.7. Attaching a cross-cap to $\mathbb{RP}^2$ (in its representation as above) produces a Klein bottle. (Exercise)

Example 8.8. Attaching a cross-handle to S^2 also produces a Klein bottle. (Exercise)

Definition 8.9. Let S_1 and S_2 be two triangulated surfaces. Choose combinatorial disks D_1 and D_2 in S_1 and S_2, respectively, that have the same number of boundary edges. Remove the interiors of D_1 and D_2 from S_1 and S_2, respectively, and glue the resulting surfaces with boundaries along the boundaries. (When we identify the two boundary cycles there are two possibilities, depending on in which direction we transverse the cycle and we have to choose one.) The result is called the *connected sum* of S_1 and S_2 and denoted by $S_1 \# S_2$.

Using Proposition 8.3 again we see that the connected sum does not depend on the chosen disks (up to combinatorial homeomorphism), only possibly on the chosen orientations of the boundaries of the disks. (Although a posteriori one sees that it does not for surfaces though it does in higher dimensions.)

Example 8.10. The connected sum of a surface S with a sphere is combinatorially homeomorphic to S again. Indeed, taking the sphere S^2 to be a tetrahedron and the disk in S to be triangle, we see that $S \# S^2$ is just subdividing the chosen triangle.

Example 8.11. Attaching a cross-cap is the same (up to combinatorial homeomorphism) as connected sum with $\mathbb{RP}^2$. Indeed, $\mathbb{RP}^2$ is S^2 with one cross-cap attached.

Likewise, the examples above show that attaching a cross-handle is the same as connected sum with the Klein bottle. Attaching a handle is the same as attaching a torus. (Exercise)

9. The conclusion

The rest is now not very difficult. We will use the result Corollary 8.4 freely.

Lemma 9.1. *Attaching two cross-caps gives a combinatorially homeomorphic result to attaching a cross-handle.*

Proof. This follows directly from Example 8.7 and Example 8.11. □

Lemma 9.2. *If the surface is not orientable, attaching a (cross-)handle does not depend up to combinatorial homeomorphism on the orientations of the triangle. This is in particular true after attaching a cross-cap.*

Proof. It suffices to produce a combinatorial homeomorphism of the surface that sends a triangle to its oppositely oriented one. Arguing as in Proposition 8.3 we just have to choose a path of a triangles that changes the orientation, which exists as our surface is not orientable. □

Now we are ready to prove our main result, which we state again.

Theorem 9.3. *Every triangulated surface S is combinatorially homeomorphic to one of the standard surfaces. More precisely, every oriented triangulated surface of Euler characteristic $2-2g$ is combinatorially homeomorphic to S_g and every non-oriented one with Euler characteristic $2-g$ to N_g.*

Proof. By Theorem 7.1, S can be obtained by attaching (cross-)handles and cross-caps to a tetrahedron. Using the lemma above, we see that we can get S by either only attaching handles (if S is orientable) or only attaching cross-caps (if S is non-orientable). Using the formula of the Euler characteristic from Theorem 7.1 we see in particular that we obtain S_g by attaching g handles to the tetrahedron (as S_g is orientable) and N_g by attaching g cross-caps (as N_g is not orientable). As we have seen that the result of attaching handles/cross-caps is up to combinatorial homeomorphism independent of the chosen disks, we obtain that every oriented triangulated surface of Euler characteristic $2-2g$ is combinatorially homeomorphic to S_g and every non-oriented one of Euler characteristic $2-g$ to N_g. □

10. Outlook

While there were already some subleties in the case of surfaces, the case of higher-dimensional manifolds is considerably more complicated. In any higher dimension a complete classification is hopeless (in dimensions ≥ 4 it is even impossible in a precise sense as it would in partiular involve deciding whether two groups with given presentation are isomorphic and Markov has proven that there is no algorithm deciding this).

If we restrict to simply-connected manifold in contrast, there is some hope. Perelman (2002) proved Poincaré's conjecture that every simply-connected 3-dimensional compact connected manifold without boundary is homeomorphic to S^3. For this he was awarded the Fields medal and one million dollar (as the Poincare conjecture was a Millenium Prize Problem), which he both declined. Freedman (1982) could already earlier classify simply-connected 4-dimensional manifolds up homeomorphism in terms of algebraic data (and obtained a Fields medal for that). For simply-connected manifolds in dimension 5 a classification is still known, but for higher-dimensional manifold it becomes more and more complicated.

This was about classification up to homeomorphism. How about classification of differentiable manifolds up to diffeomorphism? In dimension 2 or 3 it turns out that there is no difference of classifications (Moise). In contrast, in higher dimension there can be homeomorphic differentiable manifolds that are not diffeomorphic. If we talk about differentiable structures on the sphere, these are called *exotic sphere*. Milnor showed already in 1956 the existence of exotic spheres in dimension 7 (and was awarded the Fields medal for it...) and in many dimensions complete lists of exotic spheres up to diffeomorphism are known. In contrast, in dimension 4 the existence of an exotic sphere is still open and is one of the most important and difficult problems in geometric topology.

References

[Arm79] Mark Anthony Armstrong. *Basic topology.* McGraw-Hill Book Co. (UK), Ltd., London-New York, 1979. 1, 8, 10

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002. 8
[Lee11] John M. Lee. *Introduction to topological manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2011. 2, 3, 4, 8, 10
[Mas91] William S. Massey. *A basic course in algebraic topology*, volume 127 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991. 1
[Tho92] Carsten Thomassen. The Jordan-Schönflies theorem and the classification of surfaces. *Amer. Math. Monthly*, 99(2):116–130, 1992. 1, 6
[Zee66] Eric-Christopher Zeeman. An introduction to topology: The classification theorem for surfaces. `https://www.maths.ed.ac.uk/~v1ranick/surgery/zeeman.pdf`, 1966. 1, 10