THE CLASSIFICATION OF SURFACES

ANDREW PUTMAN

1. Introduction

An enormous amount of algebraic topology was developed to help classify manifolds up to homeomorphism. This classification is easy in dimensions 0 and 1, where the only connected examples are a point, a circle $\mathbb{S}^1$, and the real line $\mathbb{R}$. The first interesting dimension is 2, i.e., surfaces. Here there are infinitely many examples, but there is an elegant and easy-to-state classification (at least in the compact case) whose origins go back to 19th century work of Möbius.

1.1. **Our goal.** In this essay, we prove the classification of surfaces. Our goal is to emphasize geometric reasoning. There is a large expository gulf between the geometric topology literature and accounts of the classification of surfaces, which are typically aimed at beginning students and involve elaborate manipulations of triangulations. We include many examples and pictures, but some of our proofs and definitions are a little informal. Making them rigorous will (hopefully) be routine to readers who are experienced with smooth manifolds.

1.2. **History and sources.** The idea of our proof goes back to Zeeman [13]. Here are other accounts geared to students earlier in their education:

- See [9] or [12] for the classical combinatorial proof. I first learned this material from [9] when I was an undergraduate.
- See [2] for a proof similar to the one we give.

There are other possible proofs of this result. One that is particularly charming is Conway's "ZIP" proof, which can be found in [4]. For a history of the classification, see [5].

1.3. **Assumed results.** To avoid getting bogged down with point-set topology and foundational results about Euclidean space,[1] we will carefully state but not prove two important results:

- The existence of triangulations of surfaces. Actually, we will use the more flexible notion of "polygonal decompositions".
- The fact that the Euler characteristic of a surface is a topological invariant. This result is very easy once the theory of homology is introduced, so we see little point in giving a combinatorial proof that uses special features of surfaces.

We will also freely use standard results about smooth manifolds, often without mentioning them explicitly.

1.4. **Outline.** In §2 we give examples of surfaces and state a first version of the classification theorem. Next, as a warm-up to the proof in §3 we discuss graphs and their Euler characteristics. We then introduce polygonal decompositions and prove some basic results about the Euler characteristic in §4, which ends with a refined version of the classification. We prove the classification in the next two sections: §5 proves the "Poincaré conjecture" characterizing the sphere, and §6 proves the rest of the classification. Finally, §7 gives some extensions and generalizations of the classification.

2. Examples of surfaces

A *surface* is a 2-dimensional manifold, possibly with boundary. Our focus will be on surfaces that are connected and *closed*, that is, compact and without boundary. This section focuses on examples.

[1] Here are thinking of results like the Jordan curve theorem or more generally the Schoenflies Theorem, which says that any simple closed curve in $\mathbb{R}^2$ bounds a disk.

2.1. **Basic examples.** The most familiar surfaces are the 2-sphere $\mathbb{S}^2$ and the 2-torus $\mathbb{T}^2 = (\mathbb{S}^1)^{\times 2}$:

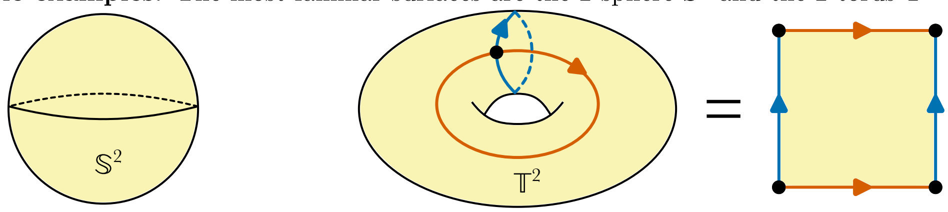

As is shown here, $\mathbb{T}^2$ can be obtained from $\mathbb{D}^2$ by identifying $\mathbb{D}^2$ with a square and identifying parallel sides. The four vertices of the square are all identified to a single point. The sphere $\mathbb{S}^2$ can also be obtained from $\mathbb{D}^2$ by identifying the entire boundary $\partial \mathbb{D}^2 = \mathbb{S}^1$ to a single point.

The torus is the surface of an ordinary donut. More generally, a *genus-g surface*, denoted Σ_g is the surface of a donut with g holes:

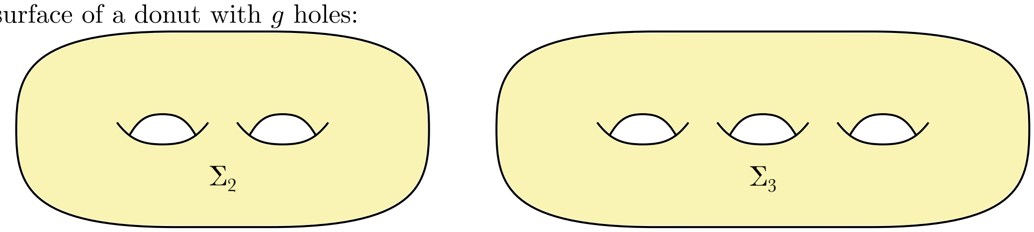

We therefore have $\Sigma_0 \cong \mathbb{S}^2$ and $\Sigma_1 \cong \mathbb{T}^2$. As we will discuss in §4 below, for $g \geq 1$ the surface Σ_g can be obtained from a $4g$-gon by identifying sides in an appropriate way.

2.2. **Möbius band and real projective plane.** The surfaces Σ_g are all orientable. We will assume that this notion is familiar from theory of manifolds. The most basic example of a non-orientable surface is a Möbius band:

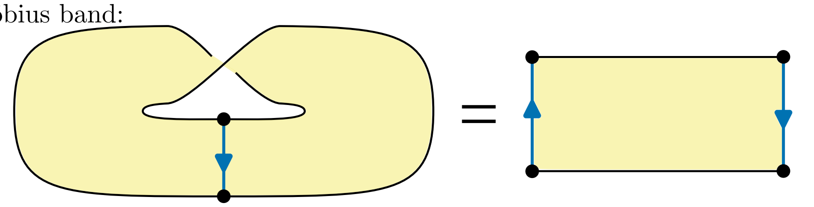

The Möbius band has one boundary component. To obtain a closed surface, we glue a disk to this boundary component to form the *real projective plane* $\mathbb{RP}^2$. Pictures of $\mathbb{RP}^2$ are not particularly enlightening,[2] but as the following shows it can be obtained from $\mathbb{D}^2$ by identifying antipodal points on the boundary $\partial \mathbb{D}^2$:

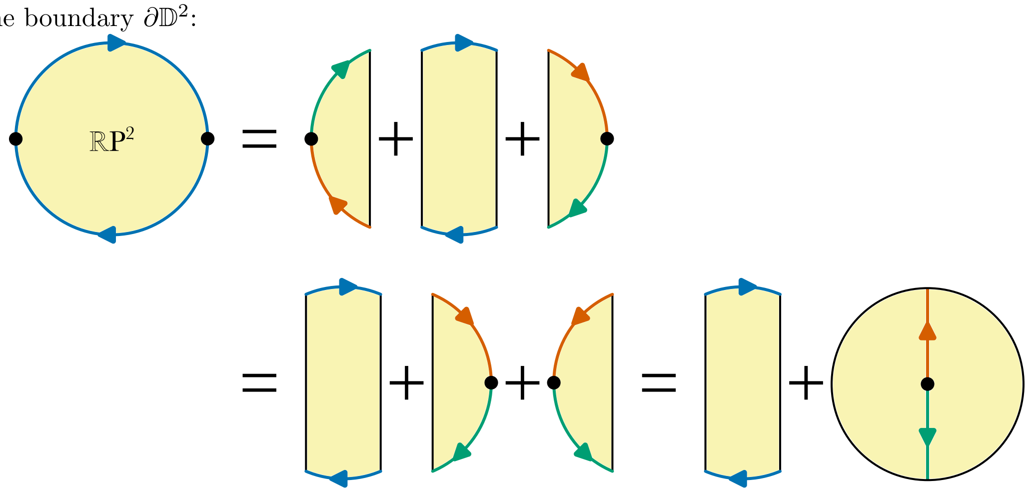

Another way of viewing $\mathbb{RP}^2$ is as the space of lines through the origin in $\mathbb{R}^3$. To connect this with the above picture, note that every such line intersects the upper hemisphere

$$U = \left\{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1,\ z \geq 0\right\} \cong \mathbb{D}^2.$$

[2]It cannot be embedded in $\mathbb{R}^3$, but only in $\mathbb{R}^4$. There is a way of drawing it in $\mathbb{R}^3$ with self-intersections called the "Boy's Surface", but this picture does not shed much light on its nature.

This intersection is unique except for lines lying in the xy-plane, which intersect

$$\partial U = \left\{(x, y, 0) \in \mathbb{R}^3 \mid x^2 + y^2 = 1\right\} \cong \partial\mathbb{D}^2$$

in two antipodal points. The space of lines through the origin can thus be identified with $U \cong \mathbb{D}^2$, but with antipodal points on $\partial U \cong \partial\mathbb{D}^2$ identified.

2.3. **Klein bottle.** Another important example of a non-orientable surface is the Klein bottle $\mathbb{K}$, which is obtained by gluing two Möbius bands together along their boundary. Unlike $\mathbb{RP}^2$, there is a somewhat enlightening way of drawing the $\mathbb{K}$, though necessarily this picture has self-intersections. See here for this and also how to get $\mathbb{K}$ by identifying the sides of a rectangle:

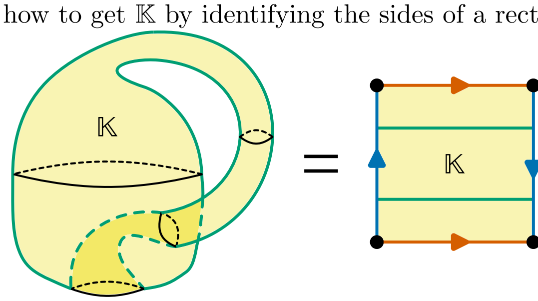

The green curve in both figures is a circle, and when you cut either open along it you get two Möbius bands. This shows that these two surfaces are indeed homeomorphic.

2.4. **Cross caps.** The surfaces $\mathbb{RP}^2$ and $\mathbb{K}$ are the first two elements of an infinite family of non-orientable surfaces. To explain this, we must introduce the notion of a *cross-cap*. A cross-cap on a surface is obtained by removing the interior of a disk and then identifying antipodal points. We denote this by drawing a disk with a cross in it like this:

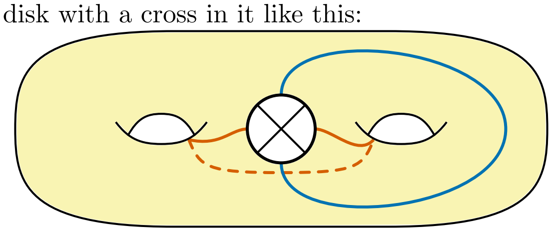

In this figure, the blue and orange arcs are actually disjoint circles embedded in the surface. As the following figure shows, a disk with a cross-cap in it is a Möbius band:

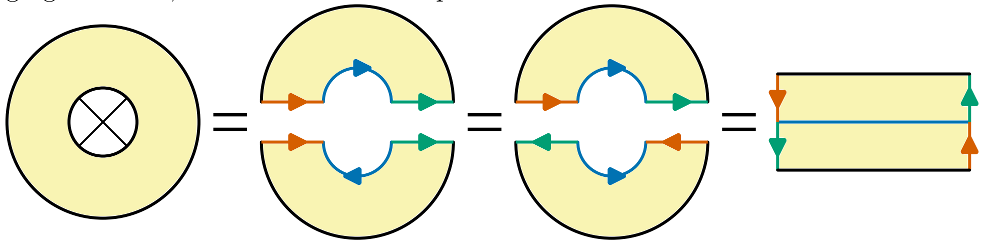

Since $\mathbb{RP}^2$ is a Möbius band with a disk glued to it and $\mathbb{K}$ is two Möbius bands glued together along their boundary, the following are $\mathbb{RP}^2$ and $\mathbb{K}$:

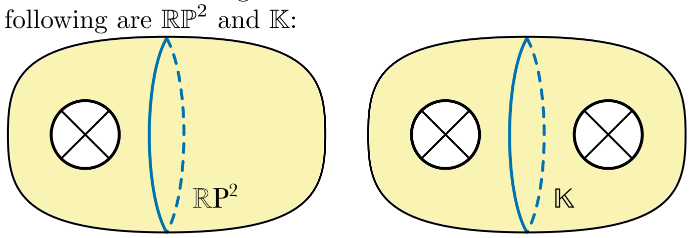

On the left the blue loop divides $\mathbb{RP}^2$ into a Möbius band and a disk, and on the right the blue loop divides $\mathbb{K}$ into two Möbius bands. These pictures suggest the general pattern: the *genus-n nonorientable surface*, denoted Σ_n^{no}, is a sphere with n cross-caps on it like this:

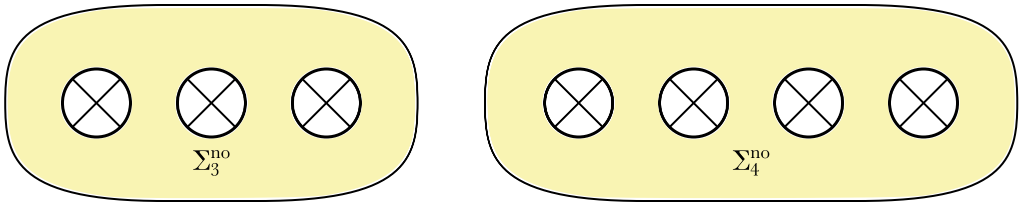

Thus $\Sigma_1^{\text{no}} \cong \mathbb{RP}^2$ and $\Sigma_2^{\text{no}} \cong \mathbb{K}$.

2.5. **Classification theorem, first version.** We can now state a first version of the classification theorem:

Theorem 2.1 (Classification of surfaces, weak)**.** *Let Σ be a closed connected surface. Then:*

- *If Σ is orientable, then $\Sigma \cong \Sigma_g$ for a unique $g \geq 0$.*
- *If Σ is non-orientable, then $\Sigma \cong \Sigma_n^{no}$ for a unique $n \geq 1$.*

In some ways this is a very satisfying result, but one weakness is that it does not give an effective way to recognize a given surface. Since it is easy to write down surfaces that do not fit into the above classification in an obvious way, this is a real problem. For instance, consider the following non-orientable surface:

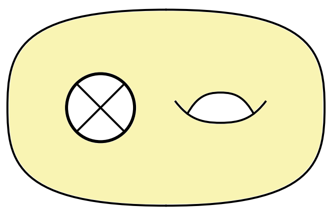

By Theorem 2.1, this must be homeomorphic to Σ_n^{no} for some $n \geq 1$. However, it is not at all obvious which Σ_n^{no} it is. We will later give a refined classification theorem that will make it clear that the above surface is Σ_3^{no}; see Theorem 4.12. Before reading this, it is worth trying to prove it for yourself.

3. Graphs and their Euler characteristics

As a warm-up before proving the classification of surfaces, this section discusses aspects of graph theory that can be viewed as a one-dimensional analogue of this classification.

3.1. **Basic definitions.** Recall that a *graph* X is a collection of vertices $V(X)$ and a collection of edges $E(X)$. Each $e \in E(X)$ connects two vertices in $V(X)$. These vertices need not be distinct:

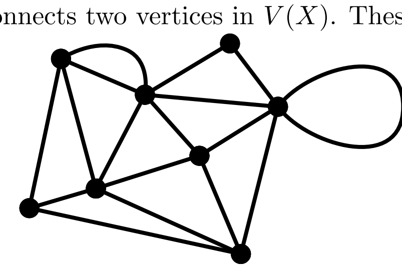

In this essay, we will only consider graphs with finitely many vertices and edges, which we call *finite graphs*. A finite graph is a topological space in a straightforward way. An *edge-path* in a graph from a vertex $v \in V(X)$ to a vertex $v' \in V(X)$ is a sequence of edges $e_1, \ldots, e_n$ such that there exist vertices

$$v = v_0, v_1, \ldots, v_n = v'$$

such that e_i connects v_{i-1} and v_i for all $1 \leq i \leq n$:

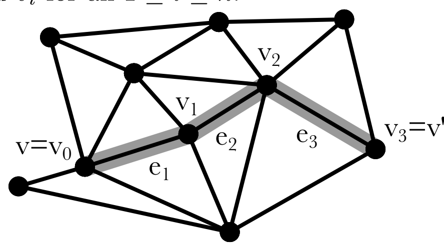

Associated to an edge-path is a continuous map $\gamma\colon I \to X$, and we will often confuse an edge-path with the associated map. The edge-path is *closed* if $v = v'$, in which case the associated path is a loop that we can regard as a continuous map $\gamma\colon \mathbb{S}^1 \to X$. The edge-path is a *cycle* if it is closed, $n \geq 1$, and all the e_i are distinct. A graph is *connected* if all distinct $v, v' \in V(X)$ are connected by an edge-path. This is equivalent to the graph being path-connected as a topological space. A *tree* is a nonempty connected graph with no cycles:

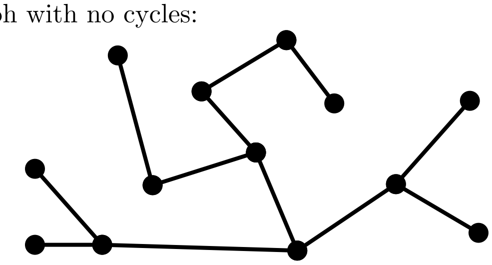

3.2. **Euler characteristic of graphs.** If X is a finite graph, then the *Euler characteristic* of X is $\chi(X) = |V(X)| - |E(X)|$.

Example 3.1. If X is the graph

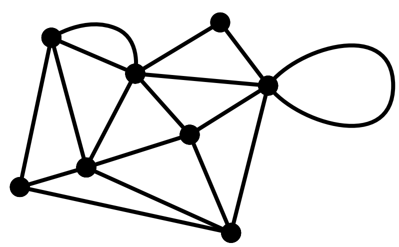

then the Euler characteristic of X is $\chi(X) = |V(X)| - |E(X)| = 8 - 17 = -9$. □

3.3. **Poincaré conjecture for graphs.** The importance of the Euler characteristic is illustrated by the following result, which we think of as the "Poincaré conjecture"[3] for graphs:

Lemma 3.2 (Poincaré conjecture for graphs)**.** *Let X be a finite nonempty connected graph. Then $\chi(X) \leq 1$, with equality if and only if X is a tree.*

Proof. If e is an edge of X connecting two distinct vertices, then we can collapse e without changing whether or not X is a tree:

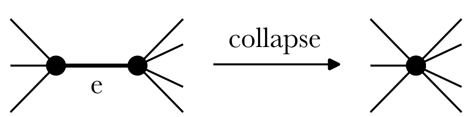

Such a collapse decreases the number of edges and vertices by 1, and thus does not change the Euler characteristic. Collapsing such edges repeatedly, we can therefore assume without loss of generality that all edges of X are loops. Since X is nonempty and connected, this implies that X has one vertex. We therefore have

$$\chi(X) = |V(X)| - |E(X)| = 1 - |E(X)| \leq 1$$

[3]For manifolds, the Poincaré conjecture is a topological characterization of a sphere. Once we have defined the Euler characteristic for surfaces, the two-dimensional Poincaré conjecture will say that a compact connected surface Σ is homeomorphic to $\mathbb{S}^2$ if and only if its Euler characteristic is 2.

with equality if and only if $|E(X)| = 0$, i.e., if and only if X is a tree.[4] □

3.4. **Maximal trees.** If X is a connected nonempty graph, then a maximal tree in X is a subtree T of X containing all the vertices. See here:

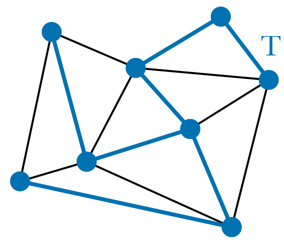

These always exist:

Lemma 3.3. *Let X be a connected nonempty graph. Then X has a maximal tree.*

This is true in general, but we only need it for finite graphs and for those the proof is slightly easier.

Proof of Lemma 3.3 for finite graphs. Assume that X has n vertices. Inductively define subtrees

$$T_1 \subset T_2 \subset \cdots \subset T_n$$

of X in the following way. Start by choosing a vertex v_0 of X and letting $T_1 = v_0$. Now assume that T_k has been constructed for some $k < n$. Since X has n vertices, it must be the case that T_k does not contain all the vertices of X. Since X is connected, this implies that there must be an edge e of X that connects a vertex of T_k to a vertex that does not lie in T_k. Let T_{k+1} be the result of adding e to T_k. This process stops at T_n, which contains n vertices and hence is a maximal tree in X. □

4. Polygonal decompositions and the Euler characteristic

Our proof of the classification of surfaces will depend on a decomposition of the surface that is a sort of two-dimensional analogue of the decomposition of a graph into vertices and edges.

4.1. **Basic definitions.** A surface equipped with a *polygonal decomposition* is a compact surface Σ (possibly with boundary) together with a finite graph X embedded in Σ such that each path component F of $\Sigma \setminus X$ is homeomorphic to an open disk $\text{Int}(\mathbb{D}^2)$. We will call such an F a *face* of the polygonal decomposition. Here is a picture of part of a polygonal decomposition, with the faces in different colors to help the reader distinguish them:

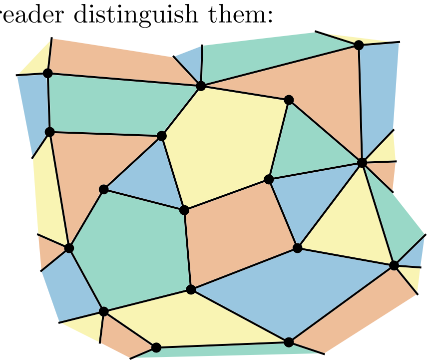

Here is some terminology for polygonal decompositions:

- The graph X will be called the 1-*skeleton*.
- The vertices and edges of X will be called the vertices and edges of the polygonal decomposition, and the sets of vertices and edges will be written $V(\Sigma)$ and $E(\Sigma)$, respectively.
- The set of faces of the polygonal decomposition will be written $F(\Sigma)$.
- The *Euler characteristic* of the polygonal decomposition is $\chi(\Sigma) = |V(\Sigma)| - |E(\Sigma)| + |F(\Sigma)|$.

[4]Since X has only one vertex, the only way it can be a tree is if it has no edges.

4.2. **Existence.** The following theorem will play a basic role in our proof:

Theorem 4.1. *Let Σ be a compact surface, possibly with boundary. Then Σ has a polygonal decomposition.*

We will not prove this theorem, which requires a long detour into point-set topology. It was originally proved by Radó [11]. See [1] and [10] for modern versions of Radó's proof. I remark that I first learned this proof from [1]. A recent and elegant proof along very different lines can be found in [6]. Amazingly, the proof in [6] uses smooth manifold techniques (even though the surface is not assumed to be smooth), and avoids doing any serious point-set work.

Remark 4.2. All of the above sources actually prove the slightly stronger fact that Σ has a triangulation, that is, a polygonal decomposition where the boundaries of each face are length-3 edge-paths in the 1-skeleton. It is easy to subdivide a general polygonal decomposition and turn it into a triangulation. □

4.3. **Examples of polygonal decompositions.** Here are a number of examples:

Example 4.3. Here are two easy polygonal decompositions of $\mathbb{S}^2$:

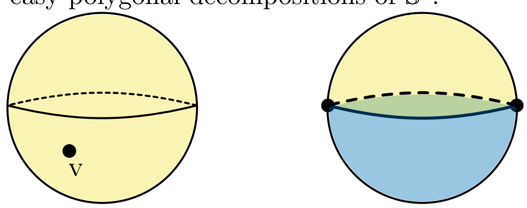

The first has a single vertex v, no edges, and one face $\mathbb{S}^2 \setminus v \cong \mathbb{R}^2 \cong \mathrm{Int}(\mathbb{D}^2)$. Its Euler characteristic is $1 - 0 + 1 = 2$. The second has two vertices, two edges, and two faces. Its Euler characteristic is $2 - 2 + 2 = 2$. Other polygonal decompositions of $\mathbb{S}^2$ can be obtained by identifying the boundaries of polyhedra in $\mathbb{R}^3$ with $\mathbb{S}^2$. For instance:

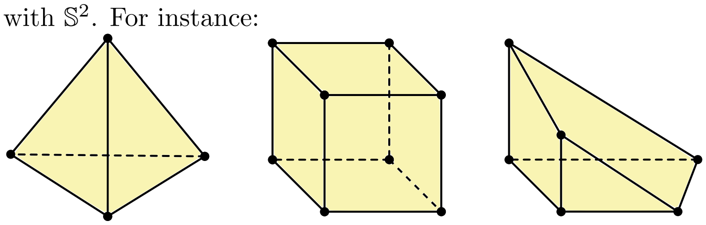

Here we have stopped trying to draw the faces in different colors. As the reader will check, in each of these cases the Euler characteristic is 2. For instance, the left-most polygonal decomposition has 4 vertices, 6 edges, and 4 faces, so its Euler characteristic is $4 - 6 + 4 = 2$. All these examples reflect a theorem we will discuss below that says that all polygonal decompositions of the same surface have the same Euler characteristic. □

Example 4.4. Here are polygonal decompositions of Σ_1 and Σ_2:

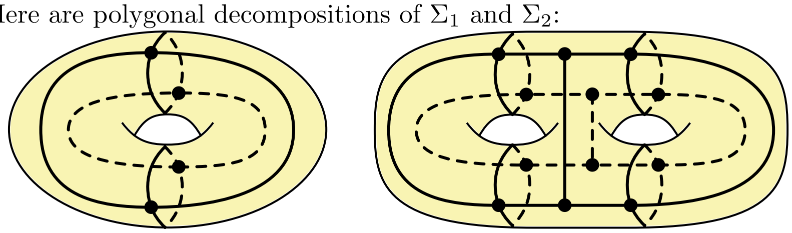

The Euler characteristics of these polygonal decompositions are

$$\chi(\Sigma_1) = 4 - 8 + 4 = 0,$$
$$\chi(\Sigma_2) = 12 - 22 + 8 = -2.$$

By regarding Σ_1 and Σ_2 as cubes with holes drilled through their centers, we can obtain two other polygonal decompositions for which it is a little easier to see that the faces are open disks:

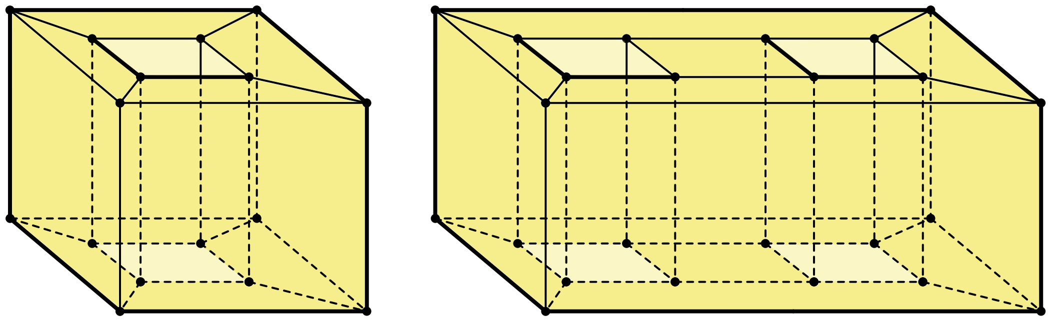

Again, the reader can verify that these have Euler characteristic 0 and -2. □

Example 4.5. When we gave examples of surfaces in §2, many of our surface came presented as polygons or disks with sides identified. This gives a polygonal decomposition of the surface. For instance, here are pictures of the torus $\mathbb{T}^2$, the real projective plane $\mathbb{RP}^2$, and the Klein bottle $\mathbb{K}$:

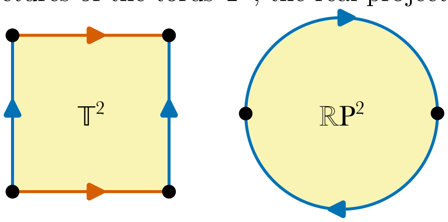

When we identify the boundary points as shown, the boundary becomes a graph and the interior of the disk/polygon gives a single face. In the above examples:

- $\mathbb{T}^2$ has one vertex and two edges and one face, so $\chi(\mathbb{T}^2) = 1 - 2 + 1 = 0$; and
- $\mathbb{RP}^2$ has one vertex and one edge and one face, so $\chi(\mathbb{RP}^2) = 1 - 1 + 1 = 1$; and
- $\mathbb{K}$ has one vertex and two edges and one face, so $\chi(\mathbb{K}) = 1 - 2 + 1 = 0$.

We will give more examples of this later in §4.10. □

4.4. **Local structure of polygonal decompositions.** Let Σ be a closed surface equipped with a polygonal decomposition. We now discuss the local structure of this polygonal decomposition. This local structure follows from the definition of a polygonal decomposition. However, the proof is a little technical and we will omit it. For a reader who cannot prove it on their own, we suggest adding these local results to the definition. The various proofs that polygonal decompositions exist (Theorem 4.1) give polygonal decompositions where this local structure definitely holds.

Our statements will be informal, but will be sufficient to understand the proof of the classification. First, a small neighborhood of a vertex looks like this:

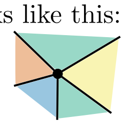

Each of the shaded regions is part of a face. These faces need not all be distinct. Next, consider an edge e. We have $e \cong [0,1]$ and $\text{Int}(e) \cong (0,1)$. A small neighborhood of $\text{Int}(e)$ looks like this:

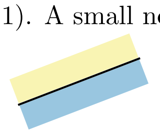

Again, the two shaded regions are parts of two faces, though these faces might be the same. We now come to a face F. Recall that $F \cong \text{Int}(\mathbb{D}^2)$. There exists a homeomorphism $\phi\colon \text{Int}(\mathbb{D}^2) \to F$ that extends to a continuous map $\phi\colon \mathbb{D}^2 \to F$. The restriction of ϕ to $\partial \mathbb{D}^2 = \mathbb{S}^1$ is usually a closed edge-path in the 1-skeleton:

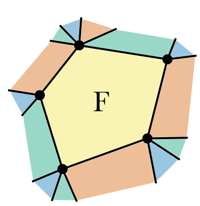

The edges in the 1-skeleton traversed by this closed edge path need not be distinct. There is one case where this does not hold: for the polygonal decomposition of $\mathbb{S}^2$ with a single vertex and face and no edges (cf. Example 4.3), the restriction of ϕ to $\partial\mathbb{D}^2 = \mathbb{S}^1$ is the constant map to this single vertex.

Remark 4.6. Polygonal decompositions also are useful for surfaces with boundary, but the local structure described above needs to be modified for vertices and edges that are contained in the boundary. □

4.5. **Well-definedness of Euler characteristic.** In our examples above, the Euler characteristics of different polygonal decompositions of the same surface were always the same. This always holds:

Theorem 4.7. *Let Σ be a compact surface, possibly with boundary. Then the Euler characteristics of any two polygonal decompositions of Σ are the same.*

The most conceptual proof of this theorem uses homology. For any reasonable compact space X, that theory produces integers $b_i(X) \geq 0$ for each $i \geq 0$ called the *Betti numbers* of X. The Betti number of X are manifestly invariants of X, and for any polygonal decomposition of a compact surface Σ we have

$$\chi(\Sigma) = b_0(\Sigma) - b_1(\Sigma) + b_2(\Sigma).$$

In fact, there are higher-dimensional versions of polygonal decompositions called *CW complex structures*. For a compact space X equipped with a CW complex structure, we have $b_i(X) = 0$ for $i \gg 0$, so the a priori infinite alternating sum

$$\chi(X) = b_0(X) - b_1(X) + b_2(X) - \cdots + (-1)^i b_i(X) + \cdots$$

is a finite sum, called the Euler characteristic of X. One of the basic structural theorems about Betti numbers is that $\chi(X)$ also equals an alternating sum of the i-dimensional cells of X.

For surfaces, there is also an alternate approach as follows. Along with the existence of polygonal decompositions (Theorem 4.1), there is also a uniqueness statement saying that any two polygonal decompositions of a compact surface Σ have what is called a *common subdivision*. What this means is that after applying a sequence of the following three moves, any two polygonal decompositions differ by a homeomorphism of Σ:

- Subdivide an edge as follows:

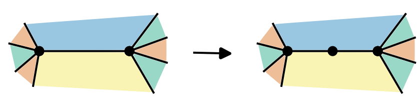

- Add a vertex and edge in the interior of a face as follows:

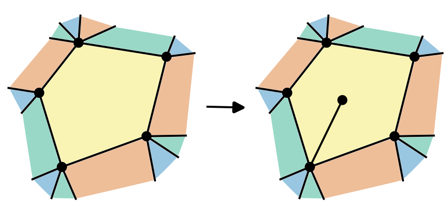

- Add an edge connecting two vertices of a face as follows:

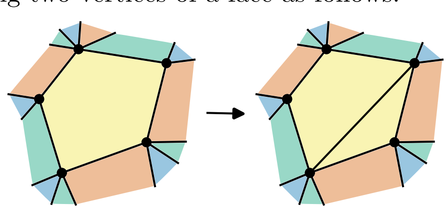

To prove that the Euler characteristic does not depend on the polygonal decomposition, it is thus enough to observe that the above three moves do not change it. For instance, subdividing an edge add one vertex and one edge, and these cancel out when calculating the Euler characteristic.

4.6. **Cutting.** We now explore the effect on the Euler characteristic of cutting along simple closed curves. Let Σ be a compact surface, possibly with boundary. Let γ be a simple closed curve lying in the interior of Σ. From the theory of manifolds, γ has what is called a *tubular neighborhood*. There are two cases:

- This tubular neighborhood is an annulus with γ in its interior like this:

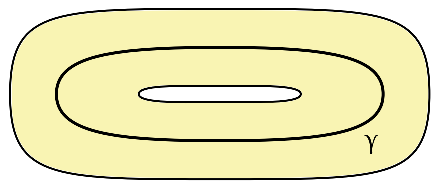

  In this case, we say that γ is a *two-sided curve*. If Σ is orientable, all γ are two-sided.
- This tubular neighborhood is a Möbius band with γ in its interior like this:

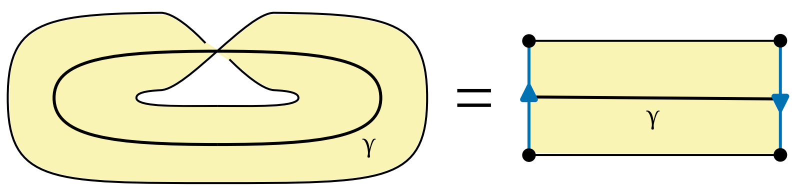

  In this case, we say that γ is a *one-sided curve*.

Cutting Σ open along γ turns Σ into a surface $\widehat{\Sigma}$. If γ is two-sided, the surface $\widehat{\Sigma}$ has two more boundary components than Σ, see here:

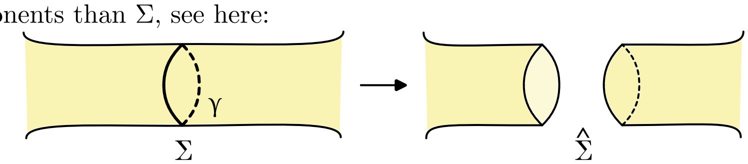

On the other hand, if γ is one-sided then $\widehat{\Sigma}$ has only more more boundary component. We remark that $\widehat{\Sigma}$ might be disconnected.

We will prove:

Lemma 4.8. *Let Σ be a compact surface, possibly with boundary. Let γ be a simple closed curve in the interior of Σ and let $\widehat{\Sigma}$ be the result of cutting Σ open along γ. Then $\chi(\widehat{\Sigma}) = \chi(\Sigma)$.*

Proof. Fix a polygonal decomposition of $\widehat{\Sigma}$. Assume first that γ is two-sided. In this case, there are boundary components ∂_1 and ∂_2 of $\widehat{\Sigma}$ such that Σ can be obtained by gluing ∂_1 to ∂_2. Necessarily ∂_1 and ∂_2 are cycles in the 1-skeleton of $\widehat{\Sigma}$. Subdividing edges in the ∂_i if necessary, we can assume that ∂_1 and ∂_2 contain the same number of n of vertices. As the following shows, Σ then has a polygonal decomposition with n fewer vertices and n fewer edges than $\widehat{\Sigma}$:

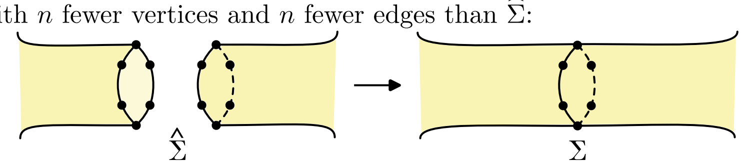

This figure does not include the part of the polygonal decomposition lying in the interior of $\widehat{\Sigma}$, which is irrelevant for this calculation. These changes in the numbers of vertices and edges cancel out when we calculate the Euler characteristic, giving that $\chi(\widehat{\Sigma}) = \chi(\Sigma)$.

Now assume that γ is one-sided. In this case, there is a boundary component ∂ of $\widehat{\Sigma}$ such that Σ can be obtained by gluing a Möbius band to ∂. Necessarily ∂ lies in the 1-skeleton of $\widehat{\Sigma}$. Assume that ∂ has n vertices. As the following shows, with respect to an appropriate polygonal decomposition Σ has 1 more edge (labeled e) and 1 more face (labeled F) than $\widehat{\Sigma}$:

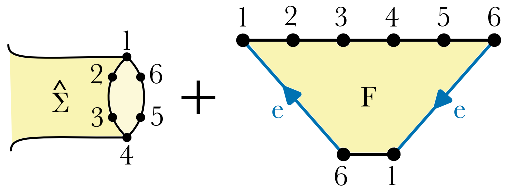

Here we have drawn the Möbius band in a skewed way to emphasize that like on ∂ its vertices are equally spaced around its single boundary component, and again we did not include the part of the polygonal decomposition lying in the interior of $\widehat{\Sigma}$. These changes in the numbers of vertices and edges again cancel out when we calculate the Euler characteristic, giving that $\chi(\widehat{\Sigma}) = \chi(\Sigma)$. □

4.7. **Capping.** We now study the effect on the Euler characteristic of gluing a disk to a boundary component:

Lemma 4.9. *Let Σ be a compact surface with boundary and let ∂ be a boundary component of Σ. Let $\overline{\Sigma}$ be the result of gluing a disk to ∂. Then $\chi(\overline{\Sigma}) = \chi(\Sigma) + 1$.*

Proof. Choose a polygonal decomposition of Σ. The disk we glued to ∂ can serve as a face of a polygonal decomposition of $\overline{\Sigma}$, giving a polygonal decomposition of $\overline{\Sigma}$ with the same number of vertices and edges as Σ and one more face. It follows that

$$\chi(\overline{\Sigma}) = |V(\overline{\Sigma})| - |E(\overline{\Sigma})| + |F(\overline{\Sigma})| = |V(\Sigma)| - |E(\Sigma)| + |F(\Sigma)| + 1 = \chi(\Sigma) + 1. \qquad \square$$

4.8. **Euler characteristic calculations.** Our results about cutting and capping make it easy to calculate the Euler characteristics of surfaces without needing to find explicit polygonal decompositions of them. As examples of this, we calculate the Euler characteristics of the genus-g surface Σ_g and the nonorientable genus-n surface Σ_n^{no}:

Lemma 4.10. *For $g \geq 0$, we have $\chi(\Sigma_g) = 2 - 2g$.*

Proof. Let $\gamma_1, \ldots, \gamma_g$ be the following two-sided curves on Σ_g:

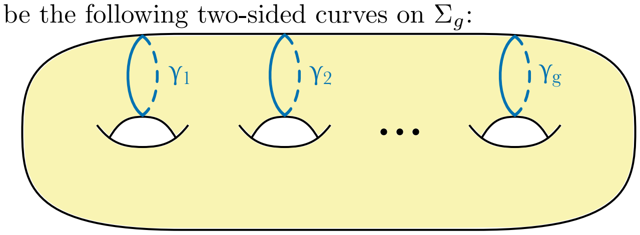

Cutting Σ_g open along the γ_i yields a connected surface $\widehat{\Sigma}$ with $2g$ boundary components. Gluing disks to each of these $2g$ boundary components yields $\mathbb{S}^2$; for instance, see here for the case $g = 3$:

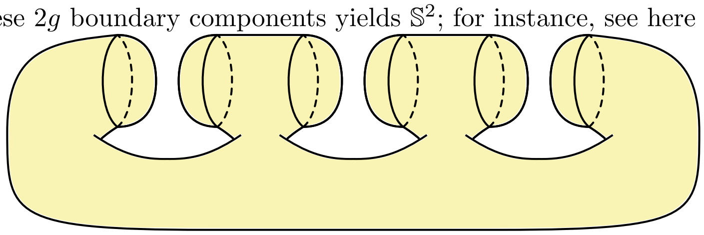

Applying Lemmas 4.8 and 4.9, we deduce that $\chi(\Sigma_g) = \chi(\widehat{\Sigma}) = \chi(\mathbb{S}^2) - 2g = 2 - 2g$. □

Lemma 4.11. *For $n \geq 1$, we have $\chi(\Sigma_n^{no}) = 2 - n$.*

Proof. Let $\gamma_1, \ldots, \gamma_n$ be the following one-sided curves on Σ_n^{no}:

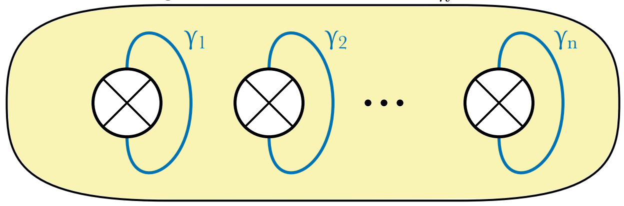

Cutting Σ_n^{no} open along the γ_i yields a connected surface $\widehat{\Sigma}$ with n boundary components. Gluing disks to each of these n boundary components yields $\mathbb{S}^2$; for instance, see here for the case $n = 2$:

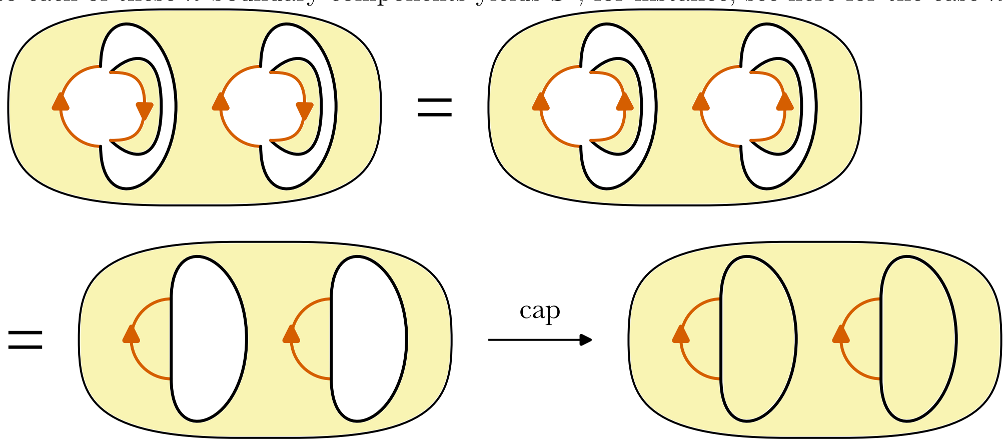

Applying Lemmas 4.8 and 4.9, we deduce that $\chi(\Sigma_n^{\text{no}}) = \chi(\widehat{\Sigma}) = \chi(\mathbb{S}^2) - n = 2 - n$. □

4.9. **Refined classification.** Recall that the classification theorem for surfaces says that every closed connected surface is homeomorphic to some Σ_g or Σ_n^{no}. Since $\chi(\Sigma_g) = 2 - 2g$ and $\chi(\Sigma_n^{\text{no}}) = 2 - n$, this will imply that the homeomorphism type of a closed connected surface is completely determined by its Euler characteristic and whether or not it is orientable. We state this refined version of the classification as follows:

Theorem 4.12 (Classification of surfaces)**.** *Let Σ be a closed connected surface. Then:*

- *If Σ is orientable then $\Sigma \cong \Sigma_g$, where $g \geq 0$ satisfies $\chi(\Sigma) = 2 - 2g$. In particular, $\chi(\Sigma)$ is even.*
- *If Σ is non-orientable then $\Sigma \cong \Sigma_n^{no}$, where $n \geq 1$ satisfies $\chi(\Sigma) = 2 - n$.*

Using this, we can answer the question we posed after stating the first version of the classification (Theorem 2.1):

Example 4.13. Consider the following surface Σ:

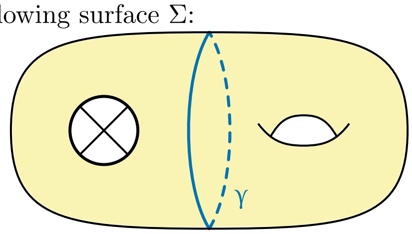

This surface is non-orientable since it contains a cross-cut. To calculate its Euler characteristic, let γ be the curve drawn above. The curve γ is two-sided, and when we cut Σ open along it and cap off the resulting two boundary components we get Σ_1 and Σ_1^{no}. It follows that

$$\chi(\Sigma) = \chi(\Sigma_1) + \chi(\Sigma_1^{\text{no}}) - 2 = 0 + 1 - 2 = -1.$$

By Theorem 4.12, we deduce that $\Sigma \cong \Sigma_3^{\text{no}}$. □

Remark 4.14. Many proofs of the classification of surfaces require the homeomorphism $\Sigma \cong \Sigma_3^{\text{no}}$ from Example 4.13, which must be proved by hand. As we will see, our proof does not need this fact, so the above argument is not circular. □

4.10. **Polygons with sides identified.** We will start the proof of Theorem 4.12 soon, but first we give a few more examples of how it can be used. One natural way to build a surface is to take a polygon (or several polygons) in $\mathbb{R}^2$ and glue its sides together. We saw several examples of this already in Example 4.5. If each side is glued to exactly one other side, then the result will always be a surface. Indeed, the only place where it is not obviously a surface is at the vertices, and a neighborhood of a vertex looks like this:

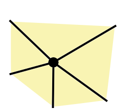

This surface has a natural polygonal decomposition where the glued-together boundary forms the 1-skeleton and the polygon is a single face (or multiple faces if there are multiple polygons). Here are some example of how to use Theorem 4.12 to identify the resulting surface:

Example 4.15. Let Σ be an octagon with sides identified as follows:

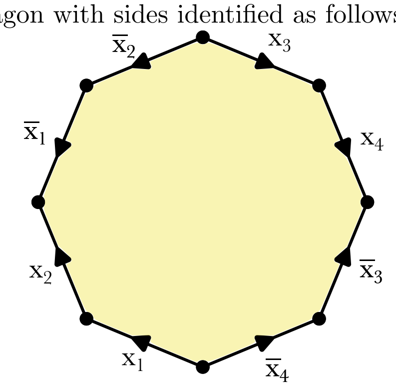

Here we have labeled the sides with letters and oriented them to show how they should be glued. The bars over the letters indicate a reversed orientation on the edge. The surface Σ is orientable since the gluing respects the orientation of the plane:

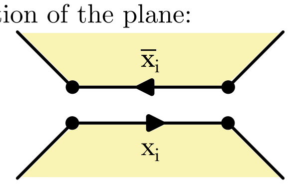

All the vertices are identified to a single vertex, and the boundary edges glue to to 4 edges. Since there is a single face, we see that $\chi(\Sigma) = 1 - 4 + 1 = -2$. It follows that $\Sigma \cong \Sigma_2$. □

Remark 4.16. A nice way to give a gluing pattern on the boundary of a $2n$-gon is to label the paired edges by letters $x_1, \ldots, x_n$. You then give a word in letters $\{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$ obtained by going around the polygon clockwise starting from some vertex and recording which edge-labels appear, with a bar indicating that the edge is traversed in the opposite orientation. For instance, in Example 4.15 the word would be $x_1 x_2 \overline{x}_1 \overline{x}_2 x_3 x_4 \overline{x}_3 \overline{x}_4$. For each $1 \leq i \leq n$, the letter x_i should appear twice (each time possibly with a bar over it). □

Example 4.17. Generalizing Example 4.15, let Σ be a $4g$-gon with sides identified according to the pattern

$$x_1 x_2 \overline{x}_1 \overline{x}_2 \cdots x_{2g-1} x_{2g} \overline{x}_{2g-1} \overline{x}_{2g}.$$

Again, all the vertices get identified to a single vertex and there are $2g$ edges and one face, so $\chi(\Sigma) = 1 - 2g + 1 = 2 - 2g$. We thus have $\Sigma \cong \Sigma_g$. □

Example 4.18. Let Σ be a $2n$-gon with sides identified according to the pattern

$$x_1 x_1 x_2 x_2 \cdots x_n x_n.$$

For instance, for $n = 3$ this is

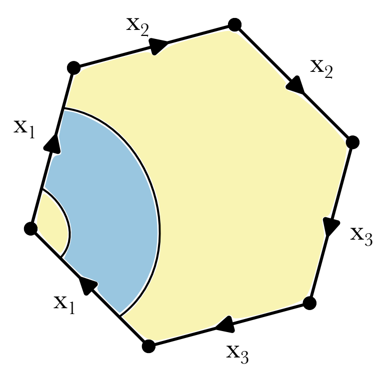

The blue strip here glues up to a Möbius band, so Σ is non-orientable. All the vertices get identified to a single vertex and there are n edges and one face, so $\chi(\Sigma) = 1 - n + 1 = 2 - n$. We thus have $\Sigma \cong \Sigma_n^{\text{no}}$. □

5. The two-dimensional Poincaré conjecture

We now start the proof of the classification by proving the two-dimensional Poincaré conjecture, which says that $\mathbb{S}^2$ is the unique connected closed surface with Euler characteristic 2:

Theorem 5.1 (Two-dimensional Poincaré conjecture)**.** *Let Σ be a closed connected surface. Then $\chi(\Sigma) \leq 2$, with equality if and only if $\Sigma \cong \mathbb{S}^2$.*

Proof. Fix a polygonal decomposition of Σ. Its 1-skeleton is a finite graph embedded in Σ. Let T be a maximal tree in the 1-skeleton. Next, let X the following graph, which one can view as a sort of "dual graph" to T:

- Put a vertex of X in the center of each polygon of our polygonal decomposition.
- Connect two vertices of X if their associated polygons share an edge that does not lie in T.

See here:

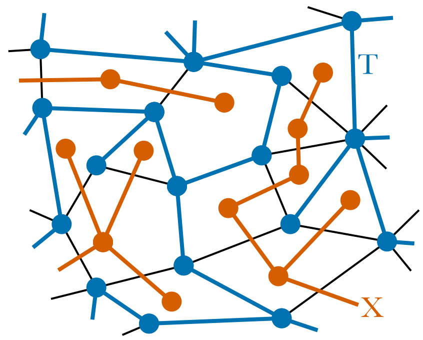

We claim that the graph X is connected. Equivalently, removing T does not disconnect Σ. The key to this is the fact that a small neighborhood U of T is homeomorphic to $\mathbb{D}^2$:

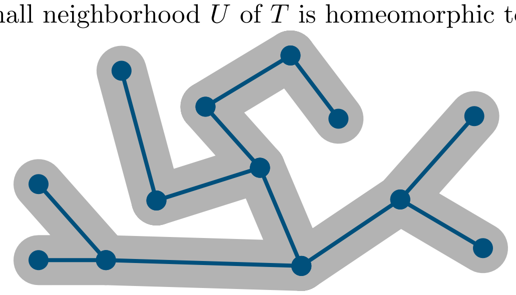

From this, we see that a path in Σ that crosses T can be re-routed to turn and follow the boundary of U until it gets to the other side of T, proving that T does not disconnect the surface.

Lemma 3.2 implies that $\chi(T) = 1$ and that $\chi(X) \leq 1$ with equality if and only if X is a tree. Since each vertex of Σ is a vertex of T, each polygon of Σ contains a unique vertex of X, and each edge of Σ is either an edge of T or is crossed by a unique edge of X, we have

$$\begin{aligned}\chi(\Sigma) &= |V(\Sigma)| - |E(\Sigma)| + |F(\sigma)| = |V(T)| - (|E(T)| + |E(X)|) + |V(X)| \\ &= \chi(T) + \chi(X) \leq 1 + 1 = 2.\end{aligned}$$

This proves half of the theorem. Equality holds if and only if $\chi(X) = 1$, i.e, if X is a tree. Assume, therefore, that X is a tree. We must prove that $\Sigma \cong \mathbb{S}^2$. To see this, note that just like above if we slightly thicken T and X we get disks D_1 and D_2. Choosing these thickenings carefully, we can ensure that D_1 and D_2 intersect in their boundaries. To help the reader understand this, here is an example of a polynomial decomposition of a surface where T and X are trees, along with D_1 and D_2:

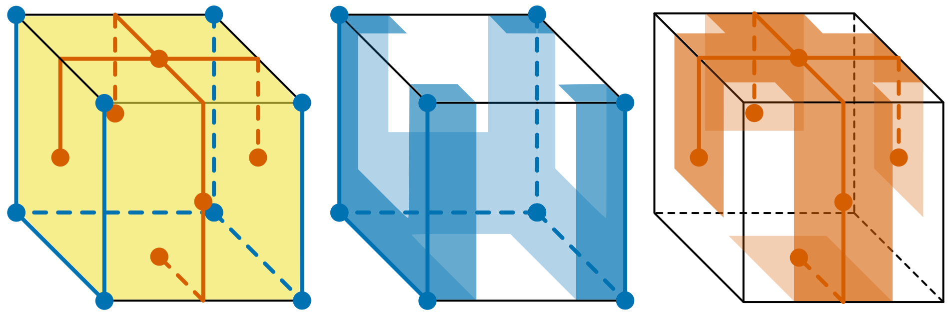

We deduce that Σ can be decomposed into two disks meeting along their boundaries:

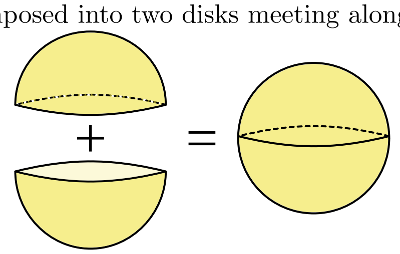

It follows that $\Sigma \cong \mathbb{S}^2$. □

Before continuing with the classification, we pause to extract the following from the above proof:

Corollary 5.2. *Let Σ be a closed connected surface such that $\chi(\Sigma) < 2$. Then there exists a simple closed curve γ on Σ that is nonseparating, i.e., such that $\Sigma \setminus \gamma$ is connected.*

Proof. Fix a polygonal decomposition of Σ, and let T and X be as in the proof of Theorem 5.1. Since $\chi(\Sigma) < 2$, it follows from the proof of Theorem 5.1 that X is a connected graph that is *not* a tree. It therefore contains a cycle γ. We claim that γ is nonseparating.

In fact, even more is true: $\Sigma \setminus X$ is connected. To see this, note that any point of $\Sigma \setminus X$ can be connected by a path in $\Sigma \setminus X$ to a vertex of the polygonal decomposition. This vertex lies in the maximal tree T, and since T is connected we can follow a path in T to any other vertex of the polygonal decomposition. The claim follows. □

6. The classification of surfaces in general

We now come to the proof of the classification of surfaces, whose statement we recall:

Theorem 4.12 (Classification of surfaces)**.** *Let Σ be a closed connected surface. Then:*

- *If Σ is orientable then $\Sigma \cong \Sigma_g$, where $g \geq 0$ satisfies $\chi(\Sigma) = 2 - 2g$. In particular, $\chi(\Sigma)$ is even.*
- *If Σ is non-orientable then $\Sigma \cong \Sigma_n^{no}$, where $n \geq 1$ satisfies $\chi(\Sigma) = 2 - n$.*

Proof. Theorem 5.1 says that $\chi(\Sigma) \leq 2$. The proof will be by reverse induction on $\chi(\Sigma)$. The base case is when $\chi(\Sigma) = 2$, in which case Theorem 5.1 says that $\Sigma \cong \mathbb{S}^2 \cong \Sigma_0$. In particular, Σ must be orientable in this case, as claimed in the theorem.

Assume now that $\chi(\Sigma) < 2$ and that the theorem is true for closed connected surfaces with larger Euler characteristics. Corollary 5.2 implies that there is a nonseparating simple closed curve γ on Σ. As we discussed in §4.6, the curve γ is either two-sided or one-sided. Let $\widehat{\Sigma}$ be the connected surface with boundary obtained by cutting Σ open along γ. There are four cases, with the first case being the only one that occurs for Σ orientable:

Case 1. *γ is two-sided and $\widehat{\Sigma}$ is orientable.*

Since γ is two-sided, $\widehat{\Sigma}$ has two boundary components. Let $\overline{\Sigma}$ be the closed connected surface obtained from $\widehat{\Sigma}$ by gluing disks to both of its boundary components. Using Lemmas 4.8 and 4.9, we have

$$\chi(\overline{\Sigma}) = \chi(\widehat{\Sigma}) + 2 = \chi(\Sigma) + 2.$$

We can therefore apply our inductive hypothesis to $\overline{\Sigma}$. Since $\widehat{\Sigma}$ is orientable, so is $\overline{\Sigma}$. It follows that $\overline{\Sigma} \cong \Sigma_g$, where $g \geq 0$ satisfies $\chi(\overline{\Sigma}) = 2 - 2g$. Since $\chi(\Sigma) = \chi(\overline{\Sigma}) - 2$, we have $\chi(\Sigma) = 2 - 2(g+1)$. Our goal, therefore, is to prove that $\Sigma \cong \Sigma_{g+1}$.

To see this, note that Σ is obtained from $\overline{\Sigma} \cong \Sigma_g$ by removing two open disks D_1 and D_2 whose closures are disjoint and gluing together the resulting boundary components. In other words, Σ is obtained by attached a handle as follows:

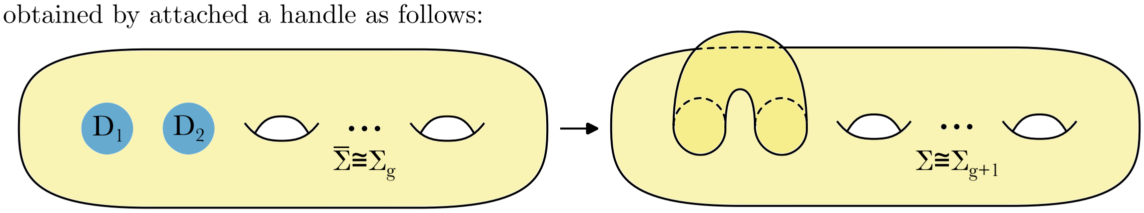

It follows that $\Sigma \cong \Sigma_{g+1}$, as desired.

Case 2. *γ is one-sided and $\widehat{\Sigma}$ is non-orientable.*

Since γ is one-sided, $\widehat{\Sigma}$ has one boundary component. Let $\overline{\Sigma}$ be the closed connected surface obtained from $\widehat{\Sigma}$ by gluing disks to its boundary component. Using Lemmas 4.8 and 4.9, we have

$$\chi(\overline{\Sigma}) = \chi(\widehat{\Sigma}) + 1 = \chi(\Sigma) + 1.$$

We can therefore apply our inductive hypothesis to $\overline{\Sigma}$. Since $\widehat{\Sigma}$ is non-orientable, so is $\overline{\Sigma}$. It follows that $\overline{\Sigma} \cong \Sigma_n^{\text{no}}$, where $n \geq 1$ satisfies $\chi(\overline{\Sigma}) = 2 - n$. Since $\chi(\Sigma) = \chi(\overline{\Sigma}) - 1$, we have $\chi(\Sigma) = 2 - (n+1)$. Our goal, therefore, is to prove that $\Sigma \cong \Sigma_{n+1}^{\text{no}}$.

To see this, note that Σ is obtained from $\overline{\Sigma} \cong \Sigma_n^{\text{no}}$ by removing an open disk D and gluing in a Möbius band. In other words, Σ is obtained by adding a cross-cap as follows:

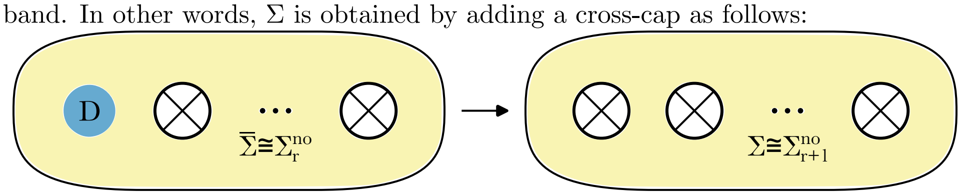

It follows that $\Sigma \cong \Sigma_{n+1}^{\text{no}}$, as desired.

Case 3. *γ is two-sided and $\widehat{\Sigma}$ is non-orientable.*

Following the argument in the previous two cases, the surface Σ has one genus and $n \geq 1$ cross-caps as follows:

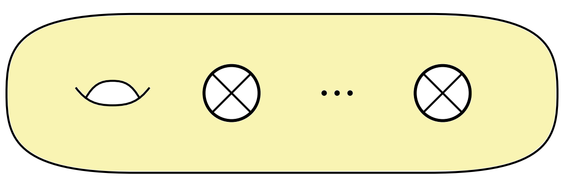

What we did wrong in this case was choose the wrong curve γ to cut along. Let γ_1 and γ_2 be as follows:

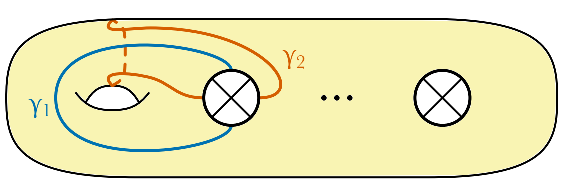

The curves γ_1 and γ_2 are both one-sided. Since γ_2 is one-sided and disjoint from γ_1, it follows that the surface $\widehat{\Sigma}'$ obtained by cutting along γ_1 is non-orientable. This reduces us to Case 2.

Case 4. *γ is one-sided and $\widehat{\Sigma}$ is orientable.*

This time, if we follow the argument from Cases 1 and 2 we get that Σ has 1 cross-cap and $g \geq 0$ genus as follows:

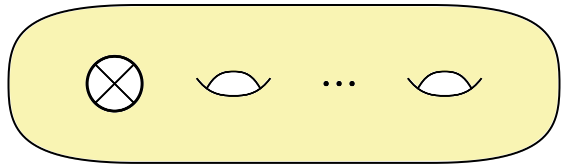

If $g = 0$ then $\Sigma \cong \Sigma_1^{\text{no}}$ and we are done. Otherwise, we can use the same trick we used in Case 3. Namely, let γ_1 and γ_2 be as follows:

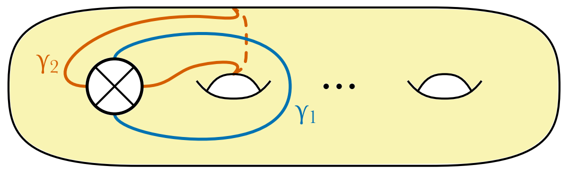

The curves γ_1 and γ_2 are both one-sided. Since γ_2 is one-sided and disjoint from γ_1, it follows that the surface $\widehat{\Sigma}'$ obtained by cutting along γ_1 is non-orientable. This reduces us to Case 2. □

7. Extensions of the classification of surfaces

We close this essay by describing two extensions of the classification of surfaces: to compact surfaces with boundary, and to non-compact surfaces.

7.1. **Compact surfaces with boundary.** Let $\Sigma_{g,b}$ be genus-g surface Σ_g with b open disks removed and let $\Sigma_{n,b}^{\text{no}}$ be a non-orientable genus-n surface with b open disks removed. For instance,

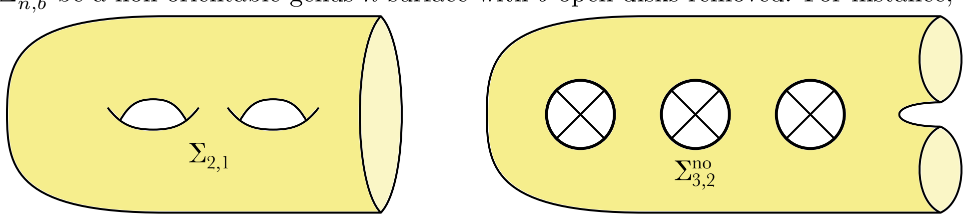

Both $\Sigma_{g,b}$ and $\Sigma_{n,b}$ are compact surfaces with b boundary components. These are the only surfaces with boundary:

Theorem 7.1 (Classification of surfaces with boundary)**.** *Let Σ be a compact connected surface with $b \geq 0$ boundary components. Then:*

- *If Σ is orientable then $\Sigma \cong \Sigma_{g,b}$, where $g \geq 0$ satisfies $\chi(\Sigma) = 2 - 2g - b$.*
- *If Σ is non-orientable then $\Sigma \cong \Sigma_{n,b}^{no}$, where $n \geq 1$ satisfies $\chi(\Sigma) = 2 - n - b$.*

Proof. Gluing disks to all the boundary components of Σ gives a closed connected surface $\overline{\Sigma}$ with $\chi(\overline{\Sigma}) = \chi(\Sigma) + b$. By the classification of surfaces (Theorem 4.12), we either have $\overline{\Sigma} \cong \Sigma_g$ with $\chi(\overline{\Sigma}) = 2 - 2g$ or $\overline{\Sigma} \cong \Sigma_n^{\text{no}}$ with $\chi(\overline{\Sigma}) = 2 - n$ depending on whether or not $\overline{\Sigma}$ (and hence Σ) is orientable. The theorem follows. □

7.2. **Noncompact surfaces.** One way to get a noncompact surface is to remove a finite number of points from the interior of a compact surface with boundary. This gives what is called a surface of *finite type*. However, non-compact surfaces can be much more complicated than this. Here are some examples.

Example 7.2. Let C be a Cantor set embedded in $\mathbb{S}^2$. Then $\mathbb{S}^2 \setminus C$ is a very complicated non-compact surface. □

Example 7.3. Consider the following infinite-genus surfaces S_1 and S_2:

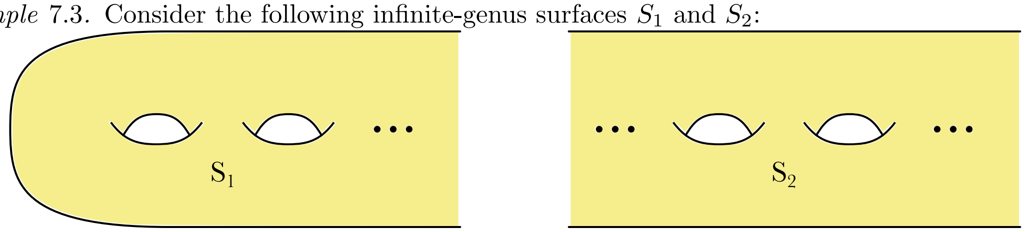

The difference between them is that S_1 has genus going off to infinity only to the right, while S_2 has genus going off to infinity in both directions. These surfaces are not homeomorphic. One way to distinguish them is to note that for every compact subset $K_1 \subset S_1$, there is only one component C of $S_1 \setminus K_1$ such that $\overline{C}$ is noncompact. However, there exist compact subsets $K_2 \subset S_2$ such that $S_2 \setminus K_2$ has two such components. This can be formalized using the theory of what are called "ends" of a space. □

This might lead the reader to think that there is no hope of classifying noncompact surfaces. However, there is such a classification making use of "end data". It was first stated by Kerékjártó [7, Chapter 5]. His proof had gaps, and the first correct proof was found by Richards [8].

Remark 7.4. Noncompact surfaces with boundary are even more complicated, especially if they have noncompact boundary components. However, a classification of them was found by Brown–Messer [3]. □

Bibliography

[1] L. V. Ahlfors and L. Sario, *Riemann surfaces*, Princeton Mathematical Series, No. 26, Princeton Univ. Press, Princeton, NJ, 1960. (Cited on page 7.)

[2] M. A. Armstrong, *Basic topology*, corrected reprint of the 1979 original, Undergraduate Texts in Mathematics, Springer, New York-Berlin, 1983. (Cited on page 1.)

[3] E. M. Brown and R. Messer, The classification of two-dimensional manifolds, Trans. Amer. Math. Soc. 255 (1979), 377–402. (Cited on page 18.)

[4] G. K. Francis and J. R. Weeks, Conway's ZIP proof, Amer. Math. Monthly 106 (1999), no. 5, 393–399 (Cited on page 1.)

[5] J. Gallier and D. Xu, A guide to the classification theorem for compact surfaces, Geometry and Computing, 9, Springer, Heidelberg, 2013. (Cited on page 1.)

[6] A. Hatcher, The Kirby torus trick for surfaces, to appear in Enseign. Math. (Cited on page 7.)

[7] B. Kerékjártó, Vorlesungen über Topologie. I, Springer, Berlin, 1923. (Cited on page 18.)

[8] I. Richards, On the classification of noncompact surfaces, Trans. Amer. Math. Soc. 106 (1963), 259–269. (Cited on page 18.)

[9] W. S. Massey, *Algebraic topology: an introduction*, reprint of the 1967 edition, Graduate Texts in Mathematics, Vol. 56, Springer, New York-Heidelberg, 1977. (Cited on page 1.)

[10] E. E. Moise, *Geometric topology in dimensions* 2 *and* 3, Graduate Texts in Mathematics, Vol. 47, Springer, New York-Heidelberg, 1977. (Cited on page 7.)

[11] T. Radó, Uber den Begriff der Riemannschen Fläche, Acta Sci. Math. Szeged. 2 (1925), 101–121. (Cited on page 7.)

[12] H. Seifert and W. Threlfall, *Seifert and Threlfall: a textbook of topology*, translated from the German edition of 1934 by Michael A. Goldman Translated from the German by Wolfgang Heil, Pure and Applied Mathematics, 89, Academic Press, New York-London, 1980. (Cited on page 1.)

[13] E. Zeeman, An introduction to topology: The classification theorem for surfaces, 1966. https://www.maths.ed.ac.uk/~v1ranick/surgery/zeeman.pdf (Cited on page 1.)

Dept of Mathematics; University of Notre Dame; 255 Hurley Hall; Notre Dame, IN 46556
Email address: andyp@nd.edu