Office Hours with a Geometric Group Theorist

Office Hours with a Geometric Group Theorist

Matt Clay and Dan Margalit, editors

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD

To...

Office Hour One

Mapping Class Groups

Tara Brendle and Leah Childers

An overarching theme in mathematics is that one can learn a vast deal about an object by studying its group of symmetries. For example, in abstract algebra we study two fundamental objects in mathematics—a finite set and a regular polygon—via the symmetric group and the dihedral group, respectively.

The primary goal of this chapter is to introduce the *mapping class group*, that is, the group of symmetries of another fundamental object: a surface. We will acquaint the reader with a few of its basic properties and give a brief glimpse of some active research related to this class of groups.

Our main goal is to find a nice generating set for the mapping class group. We will introduce elements called Dehn twists, symmetries of surfaces obtained by twisting an annulus. And we will sketch a proof of the following theorem of Max Dehn:

> *The mapping class group of a compact surface is generated by Dehn twists.*

In Section 1.1 we give an introduction to surfaces and explain the concept of a homeomorphism, our working notion of isomorphism for surfaces. In Section 1.2 we give examples of homeomorphisms and in Section 1.3 the mapping class group will be defined as a certain quotient of the group of homeomorphisms of a surface. In Section 1.4, we discuss Dehn twists and some of the relations they satisfy. We prove Dehn's theorem in Section 1.5. Finally, in Section 1.6 we list some projects and open problems.

The mapping class group is connected to many areas of mathematics, including complex analysis, dynamics, algebraic geometry, algebraic topology, geometric topology (particularly in the study of 3- and 4-dimensional spaces), and group theory. Within geometric group theory, the close relationships between mapping class groups and groups such as braid groups, Artin groups, Coxeter groups, matrix groups, and automorphism groups of free groups, have proved to be a fascinating and rich area of study. We refer you to Farb and Margalit's excellent book *A Primer on Mapping Class Groups* [4] for further details and references on many topics mentioned in this chapter. Although their text is aimed at graduate students and researchers, large portions of it are accessible to undergraduates.

1.1 A BRIEF USER'S GUIDE TO SURFACES

The word "surface" comes from the French for "on the face." Indeed, we all have an intuitive notion of a surface as the outermost layer of an object, as when we speak of resurfacing a road, or as the boundary between two substances, such as the surface of the sea. Each of these kinds of surfaces is inherently two-dimensional in nature, and mathematicians think of surfaces in similar terms.

A *homeomorphism* between surfaces (or any two topological spaces) is a continuous function with continuous inverse; equivalently, it is an invertible function f so that f and f^{-1} preserve open sets. We should think of a homeomorphism as a function that stretches and bends, but does not break or glue. For example, a circle is homeomorphic to a square since you can bend one into the other. But a circle is not homeomorphic to a line segment, since you would have to break the circle to turn it into a segment.

Exercise 1 Convince yourself that the following subsets of $\mathbb{R}^2$ are all homeomorphic: a circle, an ellipse, and a rectangle (or any polygon!). In other words, find explicit homeomorphisms between these spaces.

The official definition of a surface is: a space so that every point has an open set around it that is homeomorphic to an open set in the plane (technically, the space has to be second countable and Hausdorff, but we'll ignore this). In other words, a surface is a (second countable, Hausdorff) space where every point has an open set around it that looks planar (or a stretched, bent version of the plane). Let us give some examples.

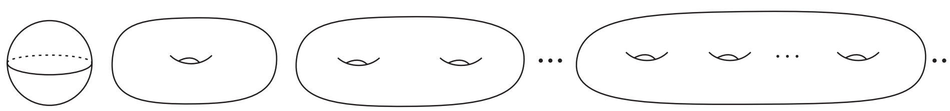

Figure 1.1 A list of surfaces.

The sphere and the torus. The leftmost surface in Figure 1.1 is familiar to us as the *sphere* S^2. We can think of S^2 as the set of points in $\mathbb{R}^3$ that are distance 1 from the origin. The next surface is the *torus* T^2, which may be familiar from calculus as a surface of revolution. For example, you can obtain a torus by taking the circle of radius 1 in the xy-plane centered at the point $(2, 0) \in \mathbb{R}^2$, and revolving it around the y-axis in $\mathbb{R}^3$.

Higher genus. The torus T^2 is often described as the frosting on a doughnut, without the doughnut. The doughnut illustration is useful for obtaining another infinite family of examples, by imagining (the frosting of) a doughnut with any number of holes, as shown in Figure 1.1. The number of holes is the *genus* of the surface. The genus can be thought of as the number of handles (or holes) on the surface.

The list of Figure 1.1 depicts surfaces increasing in genus. The sphere S^2 has genus 0. The next surface, with genus 1, is the torus T^2. The torus T^2 is followed by surfaces of genus 2 and higher.

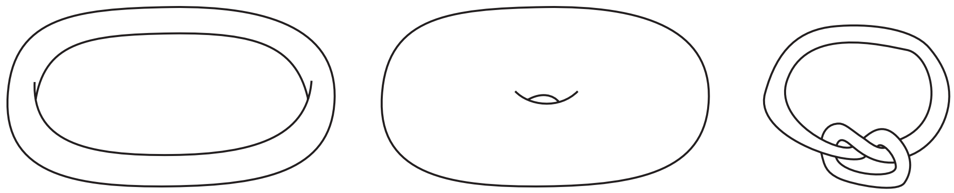

Figure 1.2 Three tori.

Three different tori? Consider the three subsets of $\mathbb{R}^3$ shown in Figure 1.2. The surface on the left is much skinnier than the torus in the middle, yet we can still recognize its basic donut shape. If we inflate the leftmost surface until it looks like the surface in the middle, we obtain a homeomorphism from the first to the second; the inverse map is obtained by deflating (notice that we can convince ourselves of this without writing down an explicit map—if you can get comfortable with this, you are becoming a topologist!). So these two surfaces, which look a little different, are really homeomorphic. So to a topoologist, they are the *same*.

But what about the other surface in Figure 1.2? We claim that it is also homeomorphic to the other two. For a moment, we imagine the first torus as a flexible hollow tube. We cut the tube, tie it in a knot, and then reglue the tube so that every point on one side of the cut is matched up exactly as before to the points on the other side. Homeomorphisms must preserve open sets, and certainly any open sets away from the cut were not disturbed by this process. But, by careful regluing, we also have not changed any of the open sets, or *neighborhoods*, of points along the circle where we cut. In fact, this process gives a homeomorphism from a standard torus to the knotted torus. The main point is that the proverbial near-sighted bug of Topology 101 cannot tell the difference between the two, because all neighborhoods remain the same—perhaps stretched a bit, but still intact. We will return to this point in the next section when we talk about homeomorphisms known as Dehn twists.

The classification of surfaces. We can think of lots of other surfaces: paraboloids, a sphere with a few points deleted, an icosahedron, a Möbius strip, the unit disk, etc. But amazingly there is a way to list them all! Let restrict to the special case of compact, orientable surfaces.

A surface is *compact* if every infinite sequence has a convergent subsequence subsequence (the surfaces shown in Figure 1.1 are compact, but the plane is not compact and a sphere or torus with finitely many points deleted is not compact).

Next, *orientable* means that we can tell the difference between clockwise and counterclockwise. A Möbius band is not orientable, because if you take a small counterclockwise loop and push at around the Möbius band, it turns into a clockwise loop. Actually, a surface is non-orientable if and only if it contains a Möbius band.

For example, a *Klein bottle*, shown in Figure 1.4 is a non-orientable surface.[1] You should try to find a Möbius strip in this surface.

[1] We're fans of the Acme Klein Bottle company; check them out at kleinbottle.com.

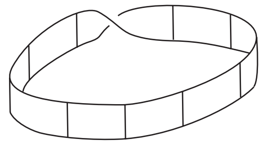

Figure 1.3 A Möbius band.

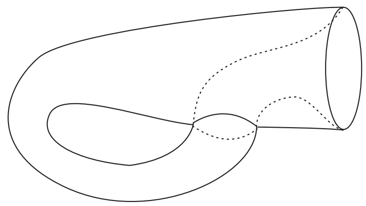

Figure 1.4 A Klein bottle.

The classification of surfaces is the following amazing fact:

> *Every compact orientable surface without boundary is homeomorphic to one of the surfaces shown in Figure 1.1.*

In other words, two compact, orientable surfaces without boundary are homeomorphic if and only if they have the same genus g, i.e. there is one compact, orientable surface without boundary for each $g \geq 0$.

From this, we can easily deduce a stronger version of the classification of surfaces:

> *Every compact orientable surface is homeomorphic to a surface obtained from one of the surfaces shown in Figure 1.1 by deleting the interiors of finitely many disjoint closed disks.*

On a first pass through this Office Hour, the student might want to ignore the case of surfaces with nonempty boundary as much as possible.

Exercise 2 Determine the genus of each of the two surfaces shown in Figure 1.5.

Exercise 3 Take a compact orientable surface S of genus one with one boundary components. The boundary of $S \times [0,1]$ is a compact surface without boundary. Which one is it? What if we start with a surface of genus g with b boundary components?

Euler characteristic. If we decompose a surface S into triangles (this means that we obtain the surface from a disjoint union of triangles by gluing edges in pairs) then the Euler characteristic $\chi(S)$ is $V - E + F$ where V, E, and F are the numbers of vertices, edges, and faces (= triangles) in the triangulation. Notice that some edges and vertices get identified in the gluing and so you need to keep track of all of this. It is an amazing fact that $\chi(S)$ does not depend on the triangulation! The Euler characteristic of a compact orientable surface of genus g with b boundary

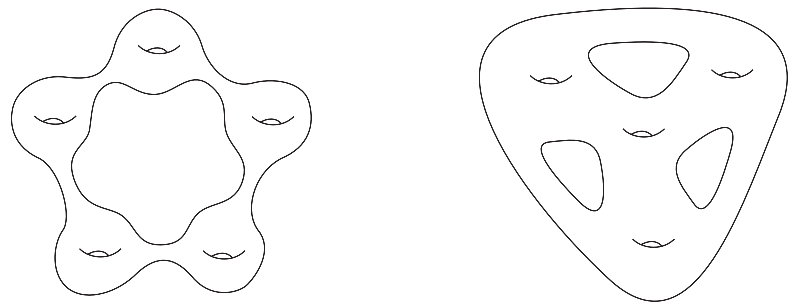

Figure 1.5 Two surfaces.

components is $2-2g-b$. It follows that a compact orientable surface is determined up to homeomorphism by any two of the three numbers χ, g, and b.

Exercise 4 Prove the last statement by finding triangulations of those surfaces (you may assume that the Euler characteristic does not depend on the triangulation).

1.2 HOMEOMORPHISMS OF SURFACES

So far we have been living in the world of topology, but the notion of homeomorphism of a surface immediately leads us to groups and group theory. Let $\mathrm{Homeo}(S)$ denote the set of homeomorphisms of S. The set $\mathrm{Homeo}(S)$ is closed under the operation of function composition. Composition is associative and by definition every homeomorphism has an inverse. We therefore see that $\mathrm{Homeo}(S)$ is a group with the identity homeomorphism as the identity element of the group.

When we first encountered homeomorphisms, we said that homeomorphic surfaces should be thought of as the same surface. In the same vein, a self-homeomorphism of a surface is precisely what we should think of as a symmetry of the surface. Normally, when we think of symmetries, we think of rigid motions, like in the dihedral group. But here in the world of topology, our symmetries are allowed to stretch and bend, but never break or glue.

The mapping class group will be defined as a quotient of (a certain subgroup of) $\mathrm{Homeo}(S)$. But before we say more about that, we introduce several important examples of elements in the group $\mathrm{Homeo}(S)$. Even though we just said that homeomorphisms are usually not rigid symmetries, the first few examples of homeomorphisms we give will in fact be rigid symmetries. These are the simplest ones to visualize.

Rotation. If we arrange our surface of genus g as in Figure 1.6, we can rotate it by $2\pi/g$, or by one "click," to obtain an element of order g in the group $\mathrm{Homeo}(S)$.

Exercise 5 Explain how the above example of a rotation of a surface of genus g gives a homeomorphism of the surface of genus g as depicted in Figure 1.1

Hyperelliptic involution. Another example of a rotation is the *hyperelliptic invo-*

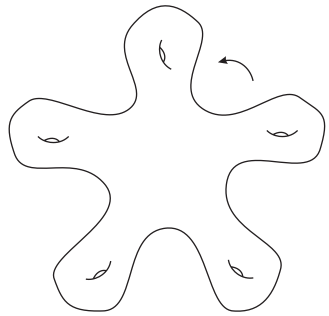

Figure 1.6 Rotation by $2\pi/g$ about the center of the surface pictured is a homeomorphism.

lution given by skewering the surface about the axis indicated in Figure 1.7 and rotating it by π.

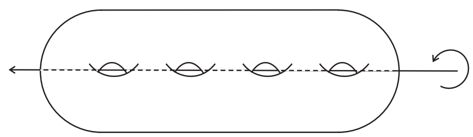

Figure 1.7 Rotation by π about the indicated axis is a hyperelliptic involution.

Reflections. Reflections of $\mathbb{R}^3$ can also give rise to homeomorphisms of a surface. As in Figure 1.8, we can just reflect across a plane that slices the surface in half. This homeomorphism is fundamentally different from the others we have discussed so far because the orientation of the surface has been reversed. If you think of writing a word on the surface, then after the reflection the words will be reversed in the same way that words look backwards in a mirror. More precisely, an orientation-reversing homeomorphism is one that takes small counterclockwise loops to small clockwise loops (we need to be on an orientable surface for this to even make sense).

Dehn twists. We come now to some homeomorphisms of surfaces that cannot be realized by rigid motions, namely, Dehn twists. Remember, one of our main goals is to convince you that these generate the mapping class group.

First, a *simple closed curve* on a surface S is the image of a circle in the surface under a continuous, injective function; three examples are shown in Figure 1.9. We can picture a simple closed curve on a surface S as a loop on the surface that does not intersect itself.

Any simple closed curve in a surface S gives rise to an important example of an element of the group $\mathrm{Homeo}(S)$. Imagine cutting a surface along a simple closed curve α, twisting one of the two resulting boundary components by a full 360-degree twist to the right, and then carefully regluing, as shown in Figure 1.10.

This is really continuous! A map is continuous when it takes nearby points to nearby points, and nearby points that are separated in the cutting are carefully reunited again when regluing (look back at the knotted torus example we discussed

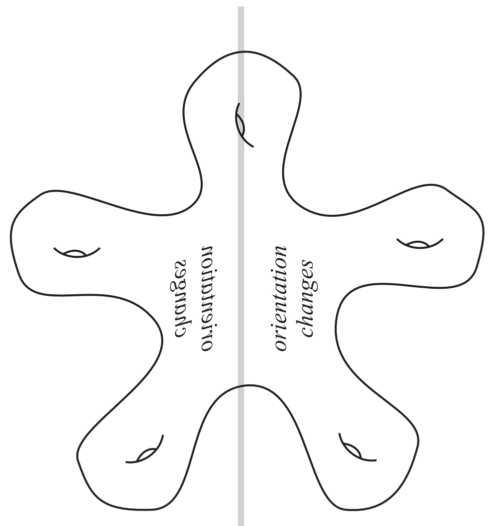

Figure 1.8 The surface is reflected across the vertical plane indicated. This homeomorphism reverses the orientation.

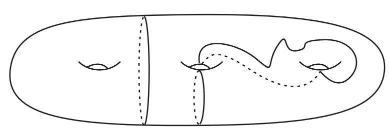

Figure 1.9 Example of three simple closed curves on a surface.

at the end of Section 1.1, which also involved cutting and regluing). The inverse is also continuous, since it is obtained by just twisting the other way. Thus we have a homeomorphism, which is known as a *Dehn twist about* α, denoted T_α.

Notice that twisting to the right makes sense as long as we have an orientation on the surface and this does not depend on any orientation of the curve we are twisting around. If we approach the curve we are twisting around from either direction, the surface gets stretched to the right. See Figure 1.11: the small horizontal arc in the left-hand picture gets twisted to the right no matter which side you approach the annulus from.

Dehn twists via annuli. We can make the definition of a Dehn twist more precise as follows. Using polar coordinates (r, θ) for points in the plane $\mathbb{R}^2$, we consider the annulus A made up of those points with $1 \leq r \leq 2$. Then we can define a map $T_A : A \to A$ by

$$(r, \theta) \mapsto (r, \theta - 2\pi r)$$

The important thing to notice is that *each point on the boundary of the annulus* A *is fixed by the map* T_A. This means that once we do our twisting on the annulus A, we can obtain an element of $\mathrm{Homeo}(S)$ by extending by the identity, that is, by fixing every other point on S outside of A. The point is that our twist on the

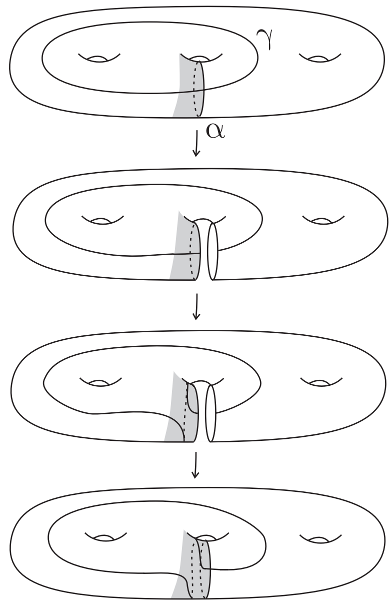

Figure 1.10 A Dehn twist seen as cut along α, twist, and reglue. The simple closed curve γ intersecting α acquires an extra twist about α.

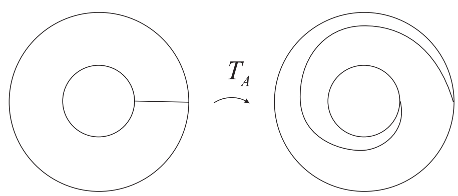

Figure 1.11 A Dehn twist on an annulus.

annulus A and the identity map on $S \backslash A$ agree where they meet, on the boundary of the annulus A.

But this discussion was supposed to be about simple closed curves, not annuli. The key realization is that every simple closed curve α in S is the core[2] of some annulus A, as in Figure 1.14 (we are only considering orientable surfaces, that is, surfaces that do not contain a Möbius band).

Exercise 6 Find a simple closed curve in the Klein bottle that is not the core of an annulus.

So given a simple closed curve α we find a corresponding annulus A, and now we're in business: we can do the Dehn twist T_A which is an element of $\mathrm{Homeo}(S)$. In fact, the map T_A we have just defined is really just T_α, a Dehn twist about the curve α as defined above. To see this, look again at Figure 1.10. We can understand

[2]In our previous discussion, the *core* of A in the plane $\mathbb{R}^2$ is the set of points with $r = \frac{3}{2}$.

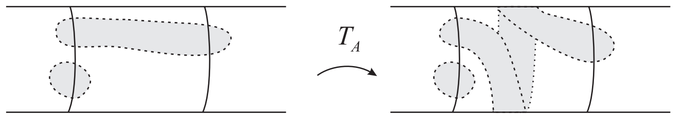

Figure 1.12 A Dehn twist preserves open sets.

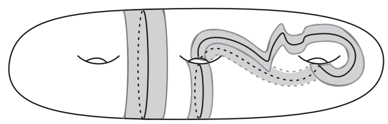

Figure 1.13 Three simple closed curves with their corresponding annuli.

this map by seeing what happens to a simple closed curve γ that crosses α: away from A, nothing happens to γ, but as it nears α, the simple closed curve γ suddenly turns and traces α before continuing on its way.

1.3 MAPPING CLASS GROUPS

In our quest to define the appropriate notion of symmetries on a surface, we have seen, through the example of Dehn twists, that the group $\mathrm{Homeo}(S)$ is somehow much too large. In order to address this, we would like to lump together homeomorphisms that are in some sense the same, and declare them to *be* the same. In other words, we are going to introduce an equivalence relation, called homotopy, on the set $\mathrm{Homeo}(S)$. The goal is to distill $\mathrm{Homeo}(S)$ into a more manageable group that still incorporates all the essential features of $\mathrm{Homeo}(S)$.

Homotopy. We like to think of homotopy as the technical tool that allows us to get away with not being very good artists when drawing simple closed curves on surfaces—a bump here or a wiggle there does not matter; drawing objects to scale is unimportant. Informally, we say two simple closed curves on a surface are *homotopic* if one can be deformed to the other; see Figure 1.15 for examples and non-examples. One way to think of this is to imagine that the simple closed curve on the surface is made of a rubber band. If you stretch the rubber band and move it around you will get a new curve homotopic to the original.

More precisely, a homotopy is a continuous deformation of one simple closed

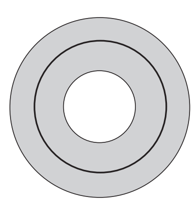

Figure 1.14 The core of an annulus.

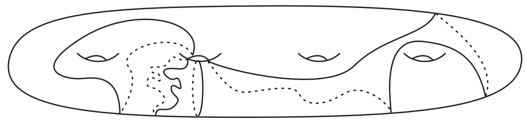

Figure 1.15 The simple closed curve on the left is not homotopic to the other three curves, which are all pairwise homotopic.

curve to another. Even more precisely, if we think of a simple closed curve in S as the image of a continuous map $S^1 \to S$, then two curves are homotopic if there is a continuous map $S^1 \times [0, 1] \to S$ so that the image of $S^1 \times \{0\}$ is the first curve and the image of $S^1 \times \{1\}$ is the second curve. We use t as the parameter for the $[0, 1]$ factor because we often think of an homotopy as a movie where at time $t = 0$ we see the first curve and then we watch the curve slowly being deformed so that by the time $t = 1$ we have arrived at the second curve.

Next we discuss homotopy for homeomorphisms. Two elements f and g in $\mathrm{Homeo}(S)$ are *homotopic* if one can be deformed to the other. More precisely, f and g are homotopic if there is a continuous map $F : S \times [0, 1] \to S$ so that $F|S \times \{0\}$ is f and $F|S \times \{1\}$ is g. Again, we can think of the homotopy F as a movie going from one homeomorphism to the other.

As a basic example (one dimension down), think of a rotation of the circle S^1 as an element of $\mathrm{Homeo}(S^1)$, and convince yourself that it is homotopic to the identity. (Even better, convince yourself that every element of $\mathrm{Homeo}(S^1)$ is homotopic to either the identity or reflection about the x-axis.)

Exercise 7 Find elements of $\mathrm{Homeo}(T^2)$ that are homotopic to the identity. Find some that are not!

Exercise 8 For any S, describe a nontrivial element of $\mathrm{Homeo}(S)$ that is homotopic to the identity.

Exercise 9 Show that "is homotopic to" is an equivalence relation on the set $\mathrm{Homeo}(S)$.

A homotopy of surface homeomorphisms is a little harder to draw and visualize than an homotopy of curves. However, it turns out that we can understand the former in terms of curves.

It is not too hard to see that if two homeomorphisms f and h of S are homotopic, then for all simple closed curves α, the curves $f(\alpha)$ and $g(\alpha)$ are homotopic. For a closed, orientable surface S of genus at least three, the converse is true:

> *If $f(\alpha)$ is homotopic to $h(\alpha)$ for all simple closed curves α then f is homotopic to h.*

This is a useful tool for showing that two homeomorphisms are homotopic as it reduces a problem about surfaces to a problem about curves. *A priori*, there are infinitely many such curves to check, but it turns out you can get away with only checking finitely many. Can you guess such a finite set of curves that determines a homeomorphism of a closed orientable surface of genus three?

If the genus is one or two, the above statement is almost true: if $f(\alpha)$ is homotopic to $h(\alpha)$ for all α then f is homotopic to either h or h times the hyperelliptic involution. For other surfaces—for instance surfaces with boundary—some version of the statement is true. Usually it is enough to consider curves and arcs instead of just curves. We will discuss the details of this as necessary.

Mapping class groups. A mapping class of a surface S is basically a homotopy class of homeomorphisms from the surface S to itself, and the mapping class group is basically the group of mapping classes of a given surface.

Our actual definition of the mapping class group will be slightly different from this. One reason for this is that our main theorem (that the mapping class group is generated by Dehn twists) will not be true otherwise, and also our inductive proof will not work otherwise.

Let S be a compact, orientable surface and let $\mathrm{Homeo}^+(S, \partial S)$ denote the group of homeomorphisms of S that preserve the orientation of S and that restrict to the identity map on each component of the boundary ∂S.

Exercise 10 Show that if S is a surface with nonempty boundary, then every homeomorphism of S that restricts to the identity on ∂S must also preserve the orientation of S.

If $h \in \mathrm{Homeo}^+(S, \partial S)$, we let $[h]$ denote the set of all homeomorphisms from S to S that are homotopic to h; here we insist that our homotopies to not move any points on the boundary of S. We say that $[h]$ is the *mapping class* of the homeomorphism h. Alternatively we say that the homeomorphism h represents the mapping class $[h]$.

The set of all mapping classes of a surface S is denoted $\mathrm{Mod}(S)$ and is called the *mapping class group* of S. Since the elements of $\mathrm{Mod}(S)$ are classes of homeomorphisms, we will use composition of homeomorphisms to define a group operation on $\mathrm{Mod}(S)$. If f and h are elements of $\mathrm{Homeo}^+(S, \partial S)$ and if $[f]$ and $[h]$ in $\mathrm{Mod}(S)$ are their respective mapping classes, then we can define an operation on $\mathrm{Mod}(S)$ as follows:

$$[f] \cdot [g] = [f \circ g].$$

It is not too hard to show that this operation is well defined and associative, that the mapping class of the identity function on S is the identity element of $\mathrm{Mod}(S)$, and that $[f]^{-1} = [f^{-1}]$ for any mapping class $[f] \in \mathrm{Mod}(S)$. Thus $\mathrm{Mod}(S)$ together with this operation is truly a group.

Exercise 11 Show that the mapping class group $\mathrm{Mod}(S)$ is the same as the quotient of $\mathrm{Homeo}^+(S, \partial S)$ by the subgroup consisting of all homeomorphisms homotopic to the identity (again, homotopies must fix the boundary pointwise).

The notation $\mathrm{Mod}(S)$ is short for *Teichmüller modular group*, an alternative name sometimes used for this group.

Dehn twists as mapping classes. Returning to our example of Dehn twists, let us consider a simple closed curve α in a surface S. Recall that in defining a Dehn twist corresponding to α, we had to make choices: an annulus A with core α, and a parametrization of the annulus A, and the resulting homeomorphism depends heavily on these choices. We seemed to have a serious problem: in the context of homeomorphisms, it makes no sense to talk about "the" Dehn twist about the simple closed curve α. Rather, we obtained an uncountably infinite number of different Dehn twists about α!

However, *the homotopy class of the resulting homeomorphism is independent of the choices*, although it is a nontrivial exercise to prove this carefully. In other words, while it does not make sense to talk about "the" Dehn twist T_α in the context of $\mathrm{Homeo}(S)$, it *does* make sense in the context of $\mathrm{Mod}(S)$.

Even better, it turns out that if α' is another simple closed curve in the surface S which is homotopic to α, then their corresponding Dehn twists are also homotopic! So not only can we choose whatever annulus and whatever parametrization we like, we are also free to choose any simple closed curve that is homotopic to α. Thus if a is a homotopy class of simple closed curves, it makes sense to write T_a as a well-defined element of $\mathrm{Mod}(S)$

Notice that—as in the case of homeomorphisms—the inverse of a Dehn twist is simply the Dehn twist about the same curve in the other direction.

We have finally arrived at the correct notion of the group of symmetries of a surface S: it is the mapping class group $\mathrm{Mod}(S)$.

1.4 DEHN TWISTS IN THE MAPPING CLASS GROUP

Recall that Dehn's theorem says that the mapping class group of a compact, orientable surface is generated by Dehn twists. To get some appreciation for this theorem let's consider the counterclockwise rotation of order 5 of the surface in Figure 1.16. By Dehn's theorem there should be a product of Dehn twists achieving this rotation. It is not at all obvious how to do this!

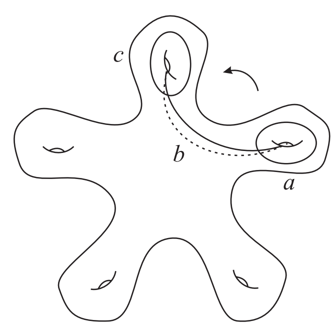

Figure 1.16 The homotopy classes a, b, and c

The rotation takes the homotopy class of curves a to the homotopy class c. As a warmup, we might simply want to find at least some product of Dehn twists that takes a to c.

First a few words about homotopy classes of curves versus curves. This is important because the mapping class group does not act on the set of curves in a surface but it does act naturally on the set of homotopy classes of curves. For two homotopy classes of curves a and b, we define the *geometric intersection number* $i(a, b)$ to be the minimum of $|\alpha \cap \beta|$ over all representatives α of a and β of b.

In order to find a product of Dehn twists taking a to c in Figure 1.16, we first observe that there is a homotopy class of curves b with $i(a, b) = i(b, c) = 1$. This is good, because we claim that:

> *If a and b are the homotopy classes of two simple closed curves that intersect in one point then $T_a T_b(a) = b$.*

Using this claim, we can take a to c by first taking a to b and then taking b to c.

So let's prove the claim. On the left-hand side of Figure 1.17 we have drawn two curves that intersect in one point; we have denoted the homotopy classes by a and b. If we multiply the desired equality $T_a T_b(a) = b$ by T_a^{-1} on both sides we obtain the equivalent equality $T_b(a) = T_a^{-1}(b)$. But this is straightforward to check; see the right-hand side of Figure 1.17.

That proves the claim! Except for one thing: it might seem like cheating that we have only checked what happens for just one pair of curves, while in fact there are infinitely many pairs of curves in a surface that intersect exactly once, even up to homotopy. The crucial point is that for every such pair of curves, there is a homeomorphism of the surface taking pair that to our pair. This homeomorphism (or rather the inverse) takes our calculation to the required calculation for the other pair. In other words, our calculation actually does *all* of the calculations! This is known as the *change of coordinates principle*; it is similar to the principle of changing basis in linear algebra.

Exercise 12 Prove that any two nonseparating simple close curves in a surface differ by a homeomorphism of the surface (a simple closed curve is nonseparating if it does not divide the surface into two pieces). *Hint: What surfaces can you get when you cut a given compact, orientable surface along a nonseparating curve?*

Exercise 13 Prove the assertion that any two pairs of simple closed curves that intersect once differ by a homeomorphism of the surface.

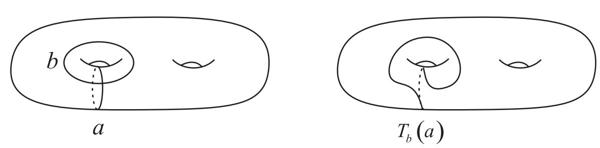

Figure 1.17 The homotopy classes a, b, and $T_a^{-1}(b) = T_b(a)$.

The problem of finding a product of Dehn twists taking one simple closed curve to another should remind you of the problem of solving a Rubik's cube using the finitely many possible twists of the Rubik's cube. In fact, there is a really fun computer game called Teruaki, written by Kazushi Ahara of Meiji University, which realizes this idea (at the time of this writing, it is available for free from his website [1]).

The braid relation. We can rephrase the claim that $T_aT_b(a) = b$ as a relation between Dehn twists in the mapping class group, namely: if $i(a,b) = 1$ then we have the relation

$$T_aT_bT_a = T_bT_aT_b.$$

This relation is called the *braid relation*.

To prove that the braid relation holds we will need the following useful fact:

> *For any $f \in \mathrm{Mod}(S)$ and any homotopy class a of simple closed curves in S we have*
>
> $$T_{f(a)} = fT_af^{-1}.$$

A bit of thought will convince you that this equation does not really require proof: following a homeomorphism to another copy of S, doing the Dehn twist there, and then going back again, is the same as if you just Dehn twist about the image of your curve under the very same homeomorphism.

Using this fact, it is easy to prove the braid relation. Indeed, the relation

$$T_aT_bT_a = T_bT_aT_b$$

is the same as

$$(T_aT_b)T_a(T_aT_b)^{-1} = T_b$$

and by our fact this is the same as

$$T_{T_aT_b(a)} = T_b$$

Now there is another fact that $T_c = T_d$ if and only if the homotopy classes c and d are the same (this is believable, but not obvious!) and so that last equality is equivalent to

$$T_aT_b(a) = b.$$

But our above claim says that this holds when $i(a,b) = 1$ so we are done!

The term "braid relation" comes from the theory of braid groups. Indeed this relation is directly connected to an analogous relation in the braid group; see Office Hour **??** for more explanation.

Groups generated by two Dehn twists. It turns out that we can completely characterize the subgroup of $\mathrm{Mod}(S)$ generated by two Dehn twists T_a and T_b in terms of $i(a,b)$. Here are the groups we get:

$i(a,b)$	$\langle T_a, T_b\rangle$
0	$\langle T_a, T_b \mid T_aT_b = T_bT_a\rangle$
1	$\langle T_a, T_b \mid T_aT_bT_a = T_bT_aT_b\rangle$
≥ 2	$\langle T_a, T_b \mid \rangle$

The first group is isomorphic to $\mathbb{Z}^2$, the second to the braid group B_3, and the third to the free group F_2. What is more, the last isomorphism can be proved using the ping pong lemma and the action of $\langle T_a, T_b\rangle$ on the set of homotopy classes of simple closed curves in the surface. The two sets in the ping pong lemma are the sets of homotopy classes of simple closed curves c with $i(a,c) > i(b,c)$ and vice versa. See the book by Farb and Margalit [4, Chapter 3] for the proofs.

1.5 GENERATING THE MAPPING CLASS GROUP BY DEHN TWISTS

Let's dive right in now and and prove Dehn's theorem:

> *For any compact, orientable surface S, the mapping class group $\mathrm{Mod}(S)$ is generated by Dehn twists.*

Even better, Stephen Humphries shoed that for a closed, orientable surface S genus g, the group $\mathrm{Mod}(S)$ is generated by Dehn twists about the $2g+1$ simple closed curves in Figure 1.18 [5].

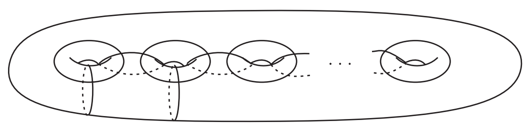

Figure 1.18 Dehn twists about these $2g+1$ simple closed curves generate the mapping class group.

As a warmup for Dehn's theorem we will convince ourselves that it is true for the compact, orientable surfaces of genus zero with 0, 1, 2, or 3 boundary components. We will later use these examples as base cases for our inductive proof of Dehn's theorem.

The disk. The first particular surface we will discuss is the disk D^2: the compact, orientable surface of genus 0 with 1 boundary component. We claim that any homeomorphism h of the disk that fixes the boundary pointwise is homotopic to the identity. Here is the homotopy: at time t, "do" h on the sub-disk of radius $1-t$ and act as the identity everywhere else (here we are taking D^2 to be the disk of radius 1 and t to vary from 0 to 1). At time 1 we just take the identity map of D^2. For each t in $[0,1]$ this gives a homeomorphism precisely because h fixes the boundary of D^2 pointwise. Further, no matter what h is, this rule defines a homotopy from h to the identity (this clever homotopy is called the Alexander trick). In particular, $\mathrm{Mod}(D^2)$ is trivial (so it is generated by Dehn twists!).

Exercise 14 Write down a precise formula in terms of h for the homotopy in the last example and verify that h is continuous.

The sphere. It is intuitively clear that any two simple closed curves in the sphere S^2 are homotopic—sketching and staring at a few pictures should convince you, although writing down a careful proof is a nontrivial exercise. It follows from this that any homeomorphism of S^2 can be modified by homotopy so that it fixes the equator pointwise (this is again intuitively clear but nontrivial to prove; there is a theorem in differential topology called the isotopy extension theorem that does the trick if you assume all of the maps in question are smooth). Any homeomorphism of S^2 that fixes the equator and preserves orientation must send the northern and southern hemispheres to themselves. But each hemisphere is just a disk, and so using the fact that $\mathrm{Mod}(D^2)$ is trivial, we conclude that our homeomorphism of S^2 is homotopic to the identity. It follows that $\mathrm{Mod}(S^2)$ is again just the trivial group.

The annulus. Let A denote the annulus $S^1 \times [0,1]$. We will argue that

$$\mathrm{Mod}(A) \cong \mathbb{Z}$$

and further that $\mathrm{Mod}(A)$ is generated by the Dehn twist about the core curve of A. The key claim is the following:

> *An arc connecting two given points on different boundary components of A is completely determined up to homotopy by how many times it winds around the circle direction of A.*

(For this claim to work we need to require that a homotopy of an arc keeps the endpoints fixed throughout the homotopy.) What exactly do we mean by the number of times that an arc winds around the circle direction of A? One way to make this precise is to choose some arc δ in A (fixed once and for all) that connects the two boundary components. Given any other arc α connecting the two boundary components we orient α so that it connects the boundary components of A in the same order as δ and then we count all of the intersections of α with δ; if α crosses δ from left to right we count a $+1$ and if it crosses from right to left we count a -1. The sum of these numbers is the desired winding number.

Let's prove the claim. Say that α and β are two arcs that both wind k times (in the same direction) around the circle direction of A. We would like to show that β is homotopic to α by a homotopy that fixes the endpoints of β. Actually, by applying the kth power of a Dehn twist about the core curve of A (or rather the inverse), we can assume without loss of generality that $k = 0$. Moreover, it doesn't hurt to assume that α is the arc δ we used in the previous paragraph to define the winding of an arc around A.

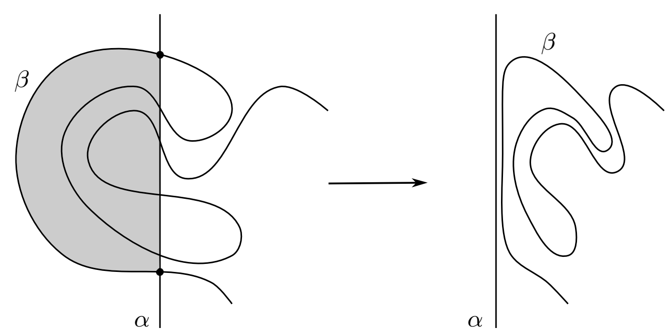

Figure 1.19 Consecutive intersection points of β with α and the disk (shaded) we use to push the arc of β

The assumption that $k = 0$ then means that β has just as many positive intersections with α as negative intersections. Therefore, if we follow along β we will find somewhere two consecutive intersections with opposite sign. The picture looks like the one in Figure 1.19; specifically, this sub-arc of β, together with the arc of α connecting the endpoints of the sub-arc of β bounds a disk in A. There may be other arcs of β intruding on this disk, but no matter: just push (or, homotope) the sub-arc of β through the disk. We have reduced the number of intersections of β with α, and continuing this process inductively we remove all intersections of α with β. The only points of intersection that remain at the end are the two points of intersection at the endpoints of α and β. But then $\alpha \cup \beta$ form a simple closed curve in A and hence bound a disk which can then be used to homotope β onto α.

Problem 1 We secretly used the Jordan–Schönflies theorem—that every simple closed curve in the plane bounds a disk—twice in the proof of the claim. Find all instances of this and explain why it is valid to apply this theorem to the annulus instead of the plane.

Now, back to showing that $\mathrm{Mod}(A)$ is isomorphic to $\mathbb{Z}$. Let α be an arc that winds 0 times around the circle direction. We can define a homomorphism to $\mathrm{Mod}(A) \to \mathbb{Z}$ whereby $f \in \mathrm{Mod}(A)$ maps to the number of times (with sign) that $f(\alpha)$ wraps around the circle direction of A. It is easy to show that this map is a surjective homomorphism: the nth power of the Dehn twist about the core of the annulus maps to n. Using the same logic we used for the sphere, we can argue it is injective. Indeed, if a homeomorphism fixes α up to homotopy, then we can modify the homeomorphism by homotopy so that it fixes α pointwise. Then there is an induced homeomorphism of the disk obtained by cutting A along α. This homeomorphism of the disk is homotopic to the identity since $\mathrm{Mod}(D^2)$ is trivial. This homotopy gives us a homotopy of the original homeomorphism of A to the identity, as desired.

The pair of pants. A *pair of pants* P is a compact, orientable surface of genus 0 with 3 boundary components (in other words, a sphere with three disks removed). We would like to show that

$$\mathrm{Mod}(P) \cong \mathbb{Z}^3$$

and moreover that $\mathrm{Mod}(P)$ is the free abelian group generated by the Dehn twists about the three boundary components of P (really, this means we take curves parallel to the boundary). Since every compact, orientable surface is made by pasting together some number of spheres, disks, annuli, and pairs of pants, this is the final ingredient we will need for our proof of Dehn's theorem.

We can take a similar tack to the one used for the annulus. Let $f \in \mathrm{Mod}(P)$. We take a simple arc α in P connecting two distinct boundary components. The key point is that an arc connecting two specific boundary components is determined up to homotopy by the number of times it winds around each boundary component. Then we can use a similar argument to the one used for the last two cases: we can modify f by Dehn twists about the two boundary components at the end points of α so that f fixes α up to homotopy; then we can cut along α to obtain an annulus, whose mapping class group we already understand.

Exercise 15 Prove that an arc in a pair of pants connecting a pair of particular points on different boundary components is determined by the number of times it winds around the boundary components at either end. You may do this by modifying the argument for the analogous claim for the annulus.

Proving Dehn's theorem. We are now ready to prove that the mapping class group of any compact, orientable surface is generated by Dehn twists. Our exposition follows closely a set of notes written by Feng Luo. [**?**] The key is the following.

Main Lemma. If c and d are simple closed curves in a compact, orientable surface S, then there is a product h of Dehn twists so that

$$i(c, h(d)) \leq 2.$$

Proof. We will show that if $i(c, d) \geq 3$ then there is a simple closed curve b so that $i(c, T_b(d)) < i(c, d)$.

The idea is to look at the pattern of intersections along c. We can draw c as a vertical arc on the page (imagine that this is a small piece of c) and draw the intersections of d with c, so that d looks like a collection of horizontal arcs. All of these arcs are connected up somewhere outside the picture, but we do not need to worry about exactly how they are connected.

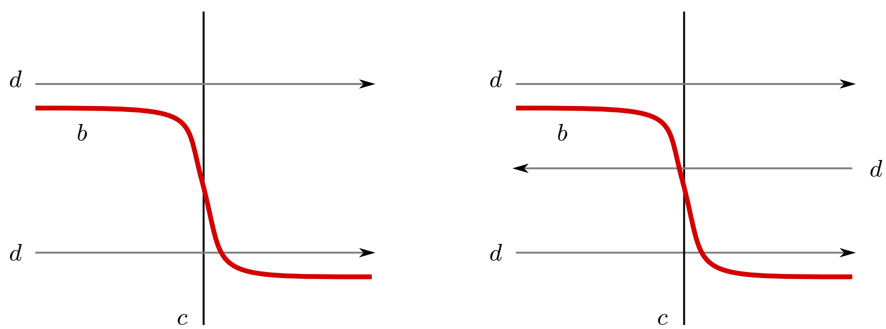

Figure 1.20 The two cases for the main lemma

We also orient c and d arbitrarily. All we care about is whether signs of intersection agree or disagree, and this does not depend at all on how we orient the two curves.

If the intersection number $i(c, d)$ is at least 3, then along the vertical arc in our picture we either have to see two consecutive intersections of the same sign or three consecutive intersections with alternating signs. In either case we can find a simple closed curve b with the desired property. The curve b is indicated in Figure 1.20. We have only shown a small part of b in the picture—after it leaves the page, b just follows along d.

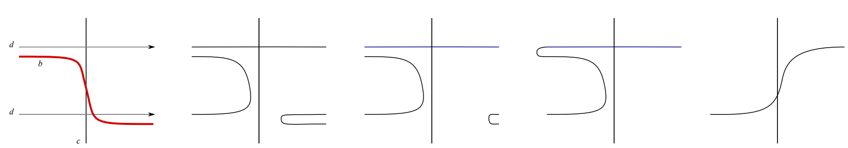

Figure 1.21 Checking the first case of the main lemma

It is a straightforward computation to check that b has the desired property. We show the computation for the first case in Figure 1.21. The first of the five pictures in the figure just shows c, d, and b, as in Figure 1.20. The second picture shows $T_b(d)$ (again, this curve follows d outside the picture). We see a portion of this curve (at the bottom right) that looks like it is backtracking, so we can simplify the picture by pushing this portion to the right as in the third picture. In the fourth picture we have pushed around d until we arrive at the top left of the picture. If we keep pushing, we get the fifth picture, at which point we see that $T_b(d)$ has (at least) one fewer point of intersection with c than d does, as desired. We leave the computation for the second case as an exercise. This completes the proof of the lemma.

Exercise 16 Check in the second case of the last proof that $i(c, T_b(d)) < i(c, d)$.

Proof of Dehn's theorem. We proceed by induction on the genus g of our (compact, orientable) surface S. The base case is $g = 0$, and to prove Dehn's theorem in this case we use induction on the number n of boundary components (induction inside induction!).

If n is 0, 1, 2, or 3, then we have a sphere, a disk, an annulus, or a pair of pants, and we already verified Dehn's theorem in those cases. Now suppose $n \geq 4$ and let $f \in \text{Mod}(S)$. Let c be a curve that cuts off a pair of pants in S (necessarily on the other side of c we have a surface of genus 0 with $n - 1$ boundary components).

By the Main Lemma there is a product of Dehn twists h so that $i(c, h \circ f(c)) \leq 2$. We claim that this implies that $h \circ f(c) = c$! The reason is simple. First of all since c is a separating curve (as are all curves in a surface of genus 0), we may only have even numbers of intersection. Next, we can check that if $i(c, d)$ is equal to 0 or 2 and $c \neq d$ then c and d surround different sets of boundary components, contradicting the fact that f and h act as the identity on the boundary of S (to see this, draw your favorite pictures of curves that intersect 0 or 2 times and then argue that all pairs of curves with those intersection numbers look like this).

Exercise 17 Verify the last sentence.

So what we have now is that there is a product of Dehn twists h so that $h \circ f(c) = c$. In other words, we can assume without loss of generality that f fixes c. And now we would like to argue by induction that f is a product of Dehn twists.

When we say that f fixes c we really mean that the homotopy class of homeomorphisms f fixes the homotopy class of curves c. It follows that we can choose a representative homeomorphism of f that fixes pointwise a representative curve of c (isotopy extension again). But then if we cut our surface along c we get two surfaces (a pair of pants and a surface of genus 0 with $n - 1$ boundary components, as above) and our representative of f induces a homeomorphism of each of these surfaces. By induction, the corresponding mapping classes are both equal to products of Dehn twists. But then it follows that the original mapping class f is a product of the same Dehn twists!

Now let $g \geq 1$. We assume by induction that every surface of genus $g-1$ (with any number of boundary components) satisfies the theorem. Let c be any nonseparating curve. The key point is that if we cut our surface along c then we get a surface of genus $g-1$ with two additional boundary components. (We can prove this using three facts: (1) the Euler characteristic of a surface of genus g with n boundary components is $2-2g-n$, (2) when we cut we create two additional boundary components, and (3) when we cut we do not change the Euler characteristic. If you do not believe this argument, you can instead take c to be a specific nonseparating curve and just check for that curve.)

Again, given any mapping class f, we can apply the Main Lemma to say that, without loss of generality, we have that $i(c, f(c))$ is either 0, 1, or 2. In any of these three cases we claim that we can find a curve b so that $i(c,b) = i(b, f(c)) = 1$. But we showed in Section 1.4 that if $i(c,b) = 1$ then there is a product of Dehn twists taking b to c (in the case $i(c, f(c)) = 1$ such a b exists but we clearly don't need it). Therefore we can modify f by a product of Dehn twists so that $f(c)$ is in fact equal to c. As in the genus 0 case, this gives us a mapping class of the surface of genus $g-1$ obtained by cutting along c. By induction on genus that mapping class is equal to a product of Dehn twists, and it follows that f is itself a product of Dehn twists, as desired. That does it!

Exercise 18 Verify the claim that if a and c are nonseparating curves with $i(a,c)$ equal to 0 or two then there is a curve b with $i(a,b) = i(b,c) = 1$. *Hint: there are three cases:* $i(a,c) = 0$*,* $i(a,c) = 2$ *with same signs of intersection, and* $i(a,c) = 2$ *with opposite signs of intersection. Draw your favorite configuration for each case and then argue that the general case looks like your configuration.*

The torus. We end with a discussion of the torus. While it follows from Dehn's theorem (which we just proved) that $\mathrm{Mod}(T^2)$ is generated by Dehn twists, we will see that this mapping class group carries the extra structure of a linear group. Indeed, we can show that

$$\mathrm{Mod}(T^2) \cong \mathrm{SL}(2, \mathbb{Z}).$$

Recall from Moon Duchin's Office Hour **??** that the fundamental group of T^2 is isomorphic to $\mathbb{Z}^2$. The idea of the isomorphism $\mathrm{Mod}(T^2) \cong \mathrm{SL}(2,\mathbb{Z})$ is that there is a homomorphism from $\mathrm{Mod}(T^2)$ to the group of automorphisms of the fundamental group of the torus:

$$\mathrm{Mod}(T^2) \to \mathrm{Aut}(\pi_1(T^2)).$$

Given an element of $\mathrm{Mod}(T^2)$ we can choose a representative that fixes our base point for $\pi_1(T^2)$, whence the action on $\pi_1(T^2)$. It turns out that for the torus this action is well defined, independent of the choice of representative homeomorphism (what special property of $\pi_1(T^2)$ are we using here?).

Since $\mathrm{Aut}(\pi_1(T^2)) \cong \mathrm{GL}(2,\mathbb{Z})$, the entire game now is to show that the map $\mathrm{Mod}(T^2) \to \mathrm{Aut}(\pi_1(T^2))$ is injective and that the image is exactly the subgroup of index two corresponding to matrices in $\mathrm{GL}(2,\mathbb{Z})$ with positive determinant. The first statement is proven using arguments similar to the ones already given (if a homeomorphism induces the trivial action then it fixes a curve, which we then cut along...). The second statement is proven by showing that the image matrix has positive determinant if and only if the homeomorphism preserves orientation.

Consider the two simple closed curves in the torus of Figure 1.22.

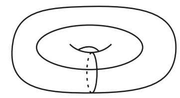

Figure 1.22 Two simple closed curves on a torus.

The Dehn twist about each of these simple closed curves is nontrivial and in fact has infinite order in $\mathrm{Mod}(T^2)$. It turns out that these two Dehn twists generate $\mathrm{Mod}(T^2)$. Indeed, it is classically known that $\mathrm{SL}(2, \mathbb{Z})$ is generated by the matrices

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}.$$

And these matrices exactly correspond to the two given Dehn twists.

Exercise 19 Complete the proof that the mapping class group of the torus is isomorphic to $\mathrm{SL}(2, \mathbb{Z})$.

1.6 PROJECTS AND OPEN QUESTIONS

As noted in the introduction, the study of $\mathrm{Mod}(S)$ is a vast area of current research involving many branches of mathematics. We will end here with a collection of projects and open problems.

Linearity. When encountering a new group, the first question we should ask is: is this group familiar? Have we ever encountered an isomorphic copy of it in some other totally different context? We just saw that the mapping class group of the torus is an example of a *linear group*, that is, it is isomorphic to a multiplicative group of matrices $\mathrm{GL}(n, \mathbb{C})$, or one of its subgroups, for some natural number n. This fact was known classically, but it was only relatively recently that Bigelow and Budney [2] gave a proof that when S is a surface of genus 2, $\mathrm{Mod}(S)$ is also a linear group, although a much more complicated one than $\mathrm{SL}(2, \mathbb{Z})$—their proof requires n to be at least 64.

There is a great deal of mathematical literature dedicated to demonstrating that mapping class groups share just about every conceivable property possessed by finitely generated linear groups, but yet the following question is currently unsolved.

Open Question 1. *Is $\mathrm{Mod}(S)$ a linear group for a surface S of any genus?*

Other generating sets. Recall that the mapping class group can be generated by $2g + 1$ Dehn twists. A reasonable question is whether we can do any better: can we generate $\mathrm{Mod}(S)$ by a smaller set of elements? It turns out we can't do better with Dehn twists, but one can generate all of $\mathrm{Mod}(S)$ with just two elements if we allow other types of elements; for instance the mapping class group is generated by two elements of finite order [8].

A discussion of various notions of "small" generating sets can be found in the introduction to [3]. As a sample, we can consider generating sets consisting of *involutions*, or elements of order 2, such as the hyperelliptic involution shown in Figure 1.7. Various mathematicians have given generating sets for $\mathrm{Mod}(S)$ consisting of involutions; for instance, Kassabov [6] and Monden [7] have shown that for certain surfaces S, just 4 involutions suffice to generate $\mathrm{Mod}(S)$. It is known that $\mathrm{Mod}(S)$ cannot be generated by two involutions (otherwise $\mathrm{Mod}(S)$ would be a quotient of the infinite dihedral group, and hence would have a cyclic group of finite index).

Open Question 2. *Can $\mathrm{Mod}(S)$ be generated by three involutions?*

Relations between higher order Dehn twists. We discussed several relations involving Dehn twists of two simple closed curves. Similar questions can be asked about the kth powers of Dehn twists. Nikolai Ivanov asked if there are "deep" relations between Dehn twists.

Open Question 3. *Is there a nontrivial relation between kth powers of Dehn twists when k is very large?*

In the last open question, the relations $T_a^k T_b^k T_a^{-k} = T^k_{T_a^k(b)}$ (and all relations that are consequences of these) are considered trivial.

Homomorphisms onto the integers. For a closed, orientable surface of genus at least 3, one can use the famous *lantern relation* in mapping class groups to prove easily that these mapping class groups are *perfect*, that is, their abelianizations are trivial [4, Chapter 5]. This implies that there are no homomorphisms from these mapping class groups onto $\mathbb{Z}$. However, this does not rule out the possibility that some finite index subgroup of the mapping class may admit such a homomorphism.

Open Question 4. *Does any finite index subgroups of the mapping class group of a closed, orientable surface of genus at least three admit a surjective homomorphism onto $\mathbb{Z}$?*

Bibliography

[1] Kazushi Ahara. Teruaki. *http://www51.atwiki.jp/kazushiahara/pages/32.html*, November 2011.

[2] Stephen Bigelow and Ryan Budney. The mapping class group of a genus two surface is linear. *Algebr. Geom. Topol*, 1(34):699–708, 2001.

[3] Tara E. Brendle and Benson Farb. Every mapping class group is generated by 6 involutions. *Journal of Algebra*, 278(1):187–198, 2004.

[4] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Princeton University Press, 2011.

[5] Stephen P. Humphries. Generators for the mapping class group. In *Topology of low-dimensional manifolds (Proc. Second Sussex Conf., Chelwood Gate, 1977)*, volume 722 of *Lecture Notes in Math.*, pages 44–47. Springer, Berlin, 1979.

[6] Martin Kassabov. Generating mapping class groups by involutions. *ArXiv Mathematics e-prints*, November 2003. arXiv:math/0311455.

[7] Naoyuki Monden. Generating the mapping class group of a punctured surface by involutions. *ArXiv e-prints*, July 2008. 0807.0916.

[8] Bronislaw Wajnryb. Mapping class group of a surface is generated by two elements. *Topology*, 35(2):377–383, 1996.

Index