Prof. Dr. Sebastian Hensel Sommersemester 2019

Mapping Class Groups

(and low-dimensional topology)

Script-in-progress

Warnings about this document

This document is a very rough script for the class; mistakes and typos may abound. If you find something problematic, please email me.

Classification of surfaces

Theorem 0.1. *Let S be a closed, oriented, connected surface.*

i) $\chi(S) \leq 2$*, with equality exactly if* $S = S^2$.

ii) If $\chi(S) < 2$*, then* S *is obtained from a surface* S' *with* $\chi(S') = \chi(S) + 2$ *by removing two disks and gluing their boundaries.*

Proof. This follows the paper "A quick proof of the classification of surfaces" by Andrew Putman (available on his website).

Choose a triangulation $\mathcal{T}$ of S. The *one-skeleton* $\mathcal{T}^1$ is the union of all vertices and edges of $\mathcal{T}$. It is an embedded graph on S. Choose a maximal tree $T \subset \mathcal{T}^1$. Such a tree contains every vertex, and does not disconnect the surface – indeed, a small ϵ–neighbourhood $U(T)$ (with respect to any chosen Riemannian metric) of the graph T is an embedded disk in S.

Construct another graph Γ, where

- Vertices v_f of Γ correspond to faces f of $\mathcal{T}$.
- Two vertices $v_f, v_{f'}$ are joined by a (geometric) edge, if the faces f, f' share an edge e *not contained in* T.

Observe that the graph Γ is connected – indeed, we will show that graphs Γ_T defined analogously for any tree T in the 1-skeleton are connected. To this end, first show that for the empty tree $T = \emptyset$, the full dual graph of $\mathcal{T}$ is clearly connected. Furthermore, if T is any tree, and e is a leaf of T, then there are paths in Γ_T connecting faces f, f' joined along an edge e in T. This follows by moving around the vertex. Now, since $T - e$ is a tree of smaller size, the claim follows by induction.

Now, let V, E, F be the numbers of vertices, edges and faces in $\mathcal{T}$. Let e_T, e_Γ be the number of edges in T, Γ. We have

$$e = e_T + e_\Gamma,$$

T has v vertices, and Γ has f vertices. Thus, we have

$$\chi(S) = v - e + f = (v - e_T) + (f - e_\Gamma) = \chi(T) + \chi(\Gamma) = 1 + \chi(\Gamma).$$

This shows that $\chi(S) \leq 2$, with equality exactly if Γ is also a tree.

In the case where Γ is a tree, we have that $U(\Gamma)$ is a disk as well, and we obtain S from joining two disks along their boundaries. This implies that S is homeomorphic to the sphere S^2.
If $\chi(S) < 2$, then Γ is not a tree, and therefore contains a loop $\alpha \subset \Gamma$. Such a loop does not disconnect S, since T is connected. Therefore, S admits a nonseparating curve, and therefore is not a sphere (by the Jordan-Schoenflies theorem). This proves i).

To prove ii), take α as above, and consider a small neighbourhood $U(\alpha)$ of α. Let $\hat{S} = S - U(\alpha)$. As S is oriented, $\hat{S}$ is a surface with two boundary components (corresponding to the two sides of α). Form a surface S' by gluing disks to the boundary components of $\hat{S}$. To compute the Euler characteristic of S', observe that in S, α was embedded in the 1–skeleton of $\mathcal{T}$. Say that α contains V_α vertices and E_α edges. This implies that we have a triangulation of $\hat{S}$ which has the same number of faces, but V_α extra vertices, and E_α extra edges. Hence, we have

$$\chi(\hat{S}) = \chi(S) + V_\alpha - E_\alpha = \chi(S)$$

as a graph which is a circle has the same number of vertices and edges. To obtain S' from $\hat{S}$ we glue two disks; each such operation has the effect of adding 1 to the Euler characteristic. This shows the theorem. $\square$

We now need the following, which is intuitive, but fairly hard to prove completely formally.

Theorem 0.2. *Suppose that S is a connected, oriented surface (possibly with boundary). Suppose that $g_i : \mathbb{D} \to S, i = 1, 2$ are two orientation preserving embeddings of disks into S. Then there is an isotopy from the identity to a map $F : S \to S$ with $F \circ g_1 = g_2$.*

Corollary 0.3. *If S is a surface, then the result of removing two disks and gluing boundaries does not depend on the choice of disks.*

Corollary 0.4. *Any two closed oriented connected surfaces with the same Euler characteristic are homeomorphic.*

Proof. Induct by Euler characteristic. The first case is from the first theorem. Otherwise, from that theorem S_1, S_2 are obtained from surfaces S_1', S_2' of the same Euler characteristic by removing and gluing disks. By induction, these are homeomorphic surfaces, and by the previous corollary the result is the same surface. $\square$

Together, these theorems imply

Theorem 0.5 (Classification of surfaces)**.** *Suppose that S, S' are two connected oriented surfaces. Then S, S' are homeomorphic if and only if S, S' have the same number of boundary components and the same Euler characteristic. If S, S' have boundary, such a homeomorphism can be chosen to induce any given permutation of the boundary components.*

Proof. Necessity is obvious. For sufficiency, let $\hat{S}, \hat{S}'$ be the surfaces obtained by gluing disks to the k boundary components to S, S'. As in the proof of the first theorem, this has the effect of adding k to the Euler characteristic. Hence, $\hat{S}, \hat{S}'$ are homeomorphic, and each have k distinguished disks corresponding to the boundary components. Inductively applying the second Theorem allows to send them to each other (in any given permutation). This yields the desired homeomorphism. $\square$

To explicitly construct the surfaces, we can e.g. glue $4g$–gons in the standard pattern and then remove some number of disks.
Another example is the following "change of coordinates principle". For its formulation, and for later, we define for a simple closed curve α the surface

$$S - \alpha = S \setminus U'(\alpha)$$

where $U(\alpha) = S^1 \times (-\epsilon, \epsilon)$ is a regular neighbourhood, and $U'(\alpha) = S^1 \times (-\epsilon/2, \epsilon/2)$. We can identify $S - \alpha$ as a subsurface of S. Also, S is obtained (up to homeomorphism) from gluing the two boundary components of $S - \alpha$.

Lemma 0.6. *Suppose S is a surface, and α, β are two simple closed curves so that $S - \alpha, S - \beta$ are both connected. Then there is a mapping class ϕ of S with $\phi(\alpha) = \beta$.*

Proof. Consider the surfaces $X = S - \alpha, Y = S - \beta$ obtained by cutting. There are distinguished boundary components $\partial_X^\pm, \partial_Y^\pm$ so that S is obtained from gluing these.
X, Y are both connected, have the same number of boundary components (two more than S), and have the same Euler characteristic. Hence, there is a homeomorphism $f : X \to Y$. Now, choose parametrisations of the boundary components. By the Alexander trick, we may assume that f respects these, and thus glues to a homeomorphism $\hat{f} : S \to S$ which, by construction, sends α to β. $\square$

Curves and Universal Covers

We will also need to know something about universal covers of surface. Namely, we have the following

Proposition 0.7. *Let S be a closed, oriented, connected surface of genus $g \geq 2$. Then the universal cover of S is homeomorphic to the hyperbolic plane, with deck group consisting of hyperbolic isometries.*

One way to prove this involves the following steps:

- Any manifold has a universal cover.
- Any surface of genus g carries a Riemannian metric of constant curvature -1.
- Any simply connected Riemannian manifold of negative curvature -1 is isometric to hyperbolic space (of the correct dimension).

Another way to do this is to build the cover explicitly, by realising that the gluing of the $4g$–gon can be extended to a tiling of the hyperbolic plane. Both of these would take us a little bit away from our main theme, but both will be (essentially) discussed in the Riemannian geometry course.

Since we do use the hyperbolic point of view later, we do need a brief crash course on hyperbolic geometry.

- We call $\mathbb{H}, \mathbb{D}$ the *upper half space and disk model* for the hyperbolic plane $\mathbb{H}^2$. The *boundary at infinity* $\partial_\infty \mathbb{H}$ is the set $\mathbb{R} \cup \{\infty\}, S^1$, respectively.
- A *geodesic* in the hyperbolic plane is a Euclidean circle or line which meets $\partial_\infty \mathbb{H}$ orthogonally.
- Between any two points there is a unique geodesic segments. Any two boundary points are joined by a unique geodesic. Given a geodesic and a point on it, there is a unique orthogonal. Any two disjoint geodesic have a unique common perpendicular.
- There is a metric on $\mathbb{H}$, and in the upper half space it computes length as
$$l(\gamma) = \int \frac{1}{\Im \gamma(t)} \|\gamma'(t)\|.$$
- An *isometry* of $\mathbb{H}^2$ is a map which acts on $\mathbb{H}$ as a linear fractional transformation
$$z \mapsto \frac{az+b}{cz+d}, \quad ad - bc = 1.$$
- Isometries preserve geodesics, induce homeomorphisms of the boundary at infinity.
- Any isometry is conjugate (by an isometry) to one of the following types:
$$z \mapsto z + 1, \quad z \mapsto \lambda z,$$
fixes a point in $\mathbb{H}^2$.
- We call the first type *parabolic*, and the second type *elliptic*.
- Parabolics have one fixed point on the boundary. Every set not containing it is attracted to the fixed point.
- Parabolics preserve no geodesic.
- Parabolics move points arbitrarily small amounts.
- Hyperbolics preserve a unique geodesic, their *axis*.
- Hyperbolics move points on the axis by a fixed amount D (called *translation length*), and all other points by more than D.
- Hyperbolics ϕ have two fixed points N, S on the boundary. They have north-south dynamics: for any $N \in U, S \in V$ there is a k so that $\phi^k(\partial_\infty \mathbb{H}^2 \setminus U) \subset V$.
- For any compact set $K \subset \mathbb{H}$, we have $\phi^k(K) \to S$.

Just from these, we can understand something about compact surfaces.

Lemma 0.8. *Suppose S is a compact surface, and suppose that $S = \mathbb{H}/\Gamma$. Then Γ consists of hyperbolic elements. No two distinct elements of Γ have axes with exactly one endpoint in common. Hyperbolic elements with the same axis have a common root.*

Proof. Deck transformations cannot fix points, so there are no elliptic elements.
Since $\mathbb{H} \to S$ is a covering and a local isometry, at every point $p \in \mathbb{H}$ there is some ϵ so that

$$B_\epsilon(p) \cap \gamma B_\epsilon(p) = \emptyset,$$

for all $\gamma \neq 1$. Since S is compact, and the condition above is Γ-invariant, we can find an ϵ which works for all p. Now, parabolics move points arbitrarily small amounts, so we are done.
Finally, suppose that there would be hyperbolics ϕ, ψ with the axes A, B that share an endpoint. Then, by conjugating ψ by a power of ϕ we find a sequence of geodesics A_n converging to A. The translation length of the corresponding elements ψ_n is constant D. Now, this violates properness of the action, as we can find infinitly many orbit points in a $2D$-ball. □

Next, we need to understand how lifts of essential simple closed curves look. We call a curve *essential*, if it does not lift to a closed curve in $\widetilde{S}$.
So, let α be such a curve, and let $\widetilde{\alpha}$ be a lift. If α is essential, then there is a deck group element g_α mapping $\widetilde{\alpha}(0)$ to $\widetilde{\alpha}(1)$. By the lemma above, g_α is hyperbolic, and therefore has an axis A_α.

Lemma 0.9. *A full lift $\widetilde{\alpha} : \mathbb{R} \to \mathbb{H}$ has the same endpoints at infinity as A_α. That means: the limits $\widetilde{\alpha}(t)$ in $\overline{\mathbb{H}^2}$ exist for $t \to \pm\infty$.*
Furthermore, these endpoints depend only on the homotopy class of α.
Finally, α is homotopic to the image of A_α.

Proof. The first statement follows since $K = \widetilde{\alpha}[0,1]$ is compact, and $\mathrm{im}\widetilde{\alpha} = \cup g_\alpha^n K$.
To show the second statement, simply observe that the image of a homotopy also lives in a compact set C.
For the final statement, choose for any $s \in [0,1]$ the orthogonal geodesic segment $o_s(t)$ to A_α parametrised by constant speed so that $o_s(0) = \widetilde{\alpha}(s)$, $o_s(1) \in A_\alpha$. Then, $g_\alpha o_0(t) = o_1(t)$, and therefore $(s,t) \to o_s(t)$ descends to the desired homotopy. □

We can use this to detect intersections *up to homotopy.*

Lemma 0.10. *Suppose that A_α, A_β intersect. Then α, β cannot be made disjoint by a homotopy.*

Proof. The endpoints of $\widetilde{\alpha}, \widetilde{\beta}$ link, since the ones of A_α, A_β do. By Schoenflies this means that any curves with these endpoints need to intersect. □

Here's a concrete example:

Corollary 0.11. *Suppose that α is a simple closed geodesic, and γ_1, γ_2 are two essential curves which meet only in a point $p \in \alpha$, and leave α on opposite sides. Then the concatenation $\gamma_1 * \gamma_2$ cannot be made disjoint from α.*

Proof. Let A be an axis of g_α, and $\widetilde{p}$ be a lift of p. Let g_1, g_2 be the deck group elements corresponding to γ_1, γ_2 lifted at p. Then, $g_i A \cap A = \emptyset$, and $g_1 A, g_2 A$ live on different sides of A. Hence, $g_2 g_1$ nests one halfspace defined by A properly into itself, and the inverse does the same with the opposite halfspace. This implies that $A_{\gamma_1\gamma_2}$ crosses A_α. We are then done by the lemma. □

We now need to know a few more things about curves.

Theorem 0.12 (Transversality)**.** *Suppose that $\alpha_1, \alpha_2 : S^1 \to S$ are any two curves. Then up to homotopy we may assume that α_1, α_2 are smooth and satisfy the following:*

- *If $\alpha_i(t) = \alpha_i(s) = p$ for some $t \neq s$, then $\alpha_i'(t), \alpha_i'(s)$ span T_pS.*
- *If $\alpha_1(t) = \alpha_2(s) = p$ for some t, s, then $\alpha_1'(t), \alpha_2'(s)$ span T_pS.*

For curves on surfaces, this is not too hard. One possibility is to replace α_i by broken geodesic arcs having the desired properties and then smooth corners. Alternatively, look up "transversality" in any textbook on differential topology.
In particular, up to homotopy, curves always have a finite number of intersections and self-intersections. We can actually detect when we see the smallest possible number fairly easily. This is the so-called *bigon criterion.*

Definition 0.13. A *bigon* is a union $a \cup b$ of two embedded arcs intersecting only in its endpoints, so that $a \cup b = \partial D$ for D an embedded disk in S.

Lemma 0.14. *If for transverse simple curves α, β on $S = \mathbb{H}^2/\Gamma$ there is no bigon $a \cup b$ formed out of arcs $a \subset \alpha, b \subset \beta$, then any two full lifts $\widetilde{\alpha}, \widetilde{\beta}$ intersect in at most one point.*

Proof. Suppose not. Then there are subarcs $\widetilde{a}, \widetilde{b}$ intersecting only in its endpoints and thus bounding a disk D (Jordan curve theorem). The interior of D may intersect finitely many other arcs of $\widetilde{a}, \widetilde{b}$ (there are finitely many intersection points on any compact arc by transversality). Thus, we may choose an *innermost* disk D whose interior is disjoint from $\widetilde{a}, \widetilde{b}$. We aim to show that $gD \cap D = \emptyset$ for all $g \in \Gamma$, from which it follows that D embeds under the covering map and yields the desired bigon.
First, consider ∂D. Here, the claim follows since otherwise D would not be innermost. Suppose now that $gD \cap D \neq \emptyset$. Then (up to switching roles), we have $gD \subset D$ (as the boundaries are disjoint, and both bound unique disks in $\mathbb{H}^2$). But, any continuous map of a disk to itself has a fixed point (Brouwer fixed point theorem), violating the fact that Γ acts freely. □

Theorem 0.15 (The bigon criterion)**.** *Two transverse simple curves α, β on $S = \mathbb{H}^2/\Gamma$ are in minimal position if and only if there is no bigon $a \cup b$ formed out of arcs $a \subset \alpha, b \subset \beta$.*

Proof. Suppose there is a bigon $a \cup b$ as in the statement. Then there is a slightly larger disk $D' \supset D$, and $a' \supset a, b' \supset b$ with a', b' intersecting in two points in D', and unlinked endpoints. Hence, in D', we can reduce the number of intersections, showing that α, β are not in minimal position.
Next, suppose that there are no bigons. Consider any two full lifts α', β' to $\mathbb{H}^2$. By the previous lemma, these are disjoint or intersect in one point.
Let $g_\alpha \in \Gamma$ be the element fixing α'. Then, intersections between α, β correspond exactly to g_α–orbits of full β–lifts intersecting α'.
If we homotope β', which full lifts intersect α' does not change (as the endpoints do not change, and thus stay linked or unlinked). As we can lift the homotopy so that it commutes with g_α, this implies that the orbits which of β'–lifts which intersected α' still do so after homotopy. Hence, the number of intersections cannot be decreased by a homotopy. $\square$

Finally, some comments on regularity:

Theorem 0.16 (Zieschang)**.** *Suppose that α, β are two simple closed curves which are homotopic. Then α, β are (ambient) isotopic.*

This is also proved by an *innermost disk argument*, but the details are messy. (For the interested, a readable account can be found in the book “Geometry and spectra of compact Riemann surfaces” by Buser). The idea is fairly simply though: if α, β are disjoint and homotopic, then they need to bound an annulus (why? Use classification of surfaces! If all components of $S - \alpha \cup \beta$ would be more complicated, then α, β could not be homotopic) and therefore are isotopic. Otherwise, lift to the universal cover. Since α, β are disjoint *up to homotopy*, endpoints cannot link. Hence, suitable lifts bound a disk, and we may choose an innermost one. Now perform a modification on that disk, reducing the number of intersections.

Theorem 0.17. *i) Any orientation preserving homotopy equivalence of S is homotopic to a diffeomorphism.*
ii) Any two homotopic homeomorphisms/diffeomorphisms are isotopic (in the same category).

For the first, there are various possibilities. One is to lift the homotopy equivalence to $\mathbb{H}^2$, look at the induced map on $\partial \mathbb{H}^2$, and construct an extension to the disk using complex analysis (e.g. by Douady-Earle). The fact that homeomorphisms are isotopic to diffeomorphisms (a consequence of i) and ii)) can be done “by hand” (see Hatcher: The Kirby Torus Trick for Surfaces).
The second part for homeomorphisms can be proved using Zieschangs theorem. Namely, successively applying Zieschang on a system of curves which cut the surface into disks we may assume that the homeomorphisms are

equal on these curves. Then apply the Alexander trick. The second part for diffeomorphisms is a consequence of the Earle-Eells theorem: components of the diffeomorphism group of a surface are contractible. A proof is not really easy.

The upshot of these things is: if we are interested in objects up to homotopy (as we always are, in the study of mapping class groups!) we may assume that curves are in minimal position, all maps are smooth, all homotopies can be changed to isotopies etc. We will often do this without mention.

Our next two big goals concern generation of the mapping class group.

Theorem 0.18 (Dehn). *The mapping class group is generated by Dehn twists.*

Theorem 0.19 (Dehn). *The mapping class group is finitely generated.*

We have seen both of these for the torus.