Math 7580 – Topics in Topology:
Riemann surfaces and their moduli
Course Notes, Spring 2023

May 15, 2023

Contents

1 Constructions of Riemann surfaces

Definition 1.1. A Riemann surface is a topological surface with an atlas of charts to $\mathbb{C}$, for which the transition functions are *holomorphic*.

Constructions.

1. Multi-valued analytic functions: e.g. $\sqrt{x(x-1)(x-2)}$ corresponds to a single-valued function $(x, y) \mapsto y$ on the Riemann surface (algebraic curve) cut out by the equation $y^2 = x(x-1)(x-2)$ in $\mathbb{C}^2$ (should projectively complete).

2. General algebraic curves

3. Riemannian metrics, in particular hyperbolic metrics.

4. Polygons, e.g. regular octagon, with opposite sides glued together. Care must be taken to define charts around the point on the surface coming from the vertices of the octagon; in this case one uses the map $z \mapsto z^{1/3}$, which is holomorphic away from 0.

5. Quotients of simply connected surfaces by properly discontinuous actions by biholomorphisms:

   - $\mathbb{C}/\Lambda$, where $\Lambda \cong \mathbb{Z}^2$ is a lattice.
   - $\mathbb{H}/\Gamma$, where Γ is discrete, torsion-free subgroup of $PSL_2(\mathbb{R})$.

1.1 Uniformization.

Theorem 1.2 (Uniformization)**.** *Every simply connected Riemann surface is biholomorphic to one of:*

(i) $\hat{\mathbb{C}}$*, the Riemann sphere*

(ii) $\mathbb{C}$

(iii) $\mathbb{H}$*, the upper half plane (which is biholomorphic to the open unit disk).*

Every Riemann surface has a universal cover as a topological surface, and the Riemann surface structure lifts upstairs to the universal cover. Hence every Riemann surface is a quotient of one the three above. The only quotient of $\hat{\mathbb{C}}$ is $\hat{\mathbb{C}}$ itself. The compact quotients of $\mathbb{C}$ are exactly the tori (the non-compact quotients are $\mathbb{C}$ and the punctured plane $\mathbb{C}^*$). All other Riemann surfaces are quotients of $\mathbb{H}$.

1.2 Automorphism groups.

Lemma 1.3 (Schwarz)**.** *Let $f : \mathbb{D} \to \mathbb{D}$ be a holomorphic with $f(0) = 0$. Then $|f'(0)| \leq 1$ and $|f(z)| \leq z$ for all $z \in \mathbb{D}$. If $|f'(0)| = 1$, then f is a rotation about 0; otherwise $|f(z)| < z$ for all $z \in \mathbb{D}$.*

Proof. Consider the function $g : \mathbb{D} \to \mathbb{C}$ given by $g(z) = f(z)/z$; the value $g(0)$ is just $f'(0)$. This g is also holomorphic. By the Maximum Modulus Principle,

$$\sup |g| \leq \frac{\sup |f|}{1} \leq 1,$$

hence $|g(0)| = |f'(0)| \leq 1$, and $|f(z)| \leq z$ for all $z \in \mathbb{D}$. Equality in the above is attained only if g is a constant function, which corresponds to f a rotation. ∎

Theorem 1.4. *The groups of biholomorphic automorphisms of each of the simply connected Riemann surfaces:*

(i) $\mathrm{Aut}(\hat{\mathbb{C}}) = PSL_2(\mathbb{C})$*, acting by Mobius transformations* $z \mapsto \frac{az+b}{cz+d}$.

(ii) $\mathrm{Aut}(\mathbb{C}) = \{z \mapsto az + b : a, b \in \mathbb{C}, a \neq 0\}$

(iii) $\mathrm{Aut}(\mathbb{H}) = PSL_2(\mathbb{R})$*, acting by Mobius transformations.*

Every automorphism of $\hat{\mathbb{C}}$ has a fixed point, so we don't get any surface quotients. For $\mathbb{C}$, the quotient Riemann surfaces are the annulus $\mathbb{C}/\mathbb{Z}$ and complex tori.

Proof of (iii). It is easy to check that every element of $PSL_2(\mathbb{R})$ gives an automorphism. This group acts transitively on $\mathbb{H}$ (the affine maps $az + b$ already do). By moving to $\mathbb{D}$ and applying the Schwarz lemma, we see the only maps fixing a particular point are "rotations". The group $PSL_2(\mathbb{R})$ contains such rotations (consider the map $f(z) = \frac{\cos\theta z + \sin\theta}{(-\sin\theta)z + \cos\theta}$, which fixes i), and hence this is the full automorphism group. ∎

In the above proof, we showed that $\mathrm{Aut}(\mathbb{H})$ acts transitively on *tangent vectors*. It follows that there is a unique Riemannian metric on $\mathbb{H}$ invariant under all biholomorphic automorphisms. This turns out to be the hyperbolic metric

$$ds^2 = \frac{dx^2 + dy^2}{y^2}$$

in the coordinates $x + iy$. And $Isom^+(\mathbb{H}) = \mathrm{Aut}(\mathbb{H}) = PSL_2(\mathbb{R})$. It follows that any Riemann surface that is a quotient of $\mathbb{H}$ inherits a hyperbolic metric from it.

2 Classification of annuli

Examples of annuli:

1. $\mathbb{C}^* := \mathbb{C} - \{0\}$
2. $\mathbb{D}^*$
3. $A_{r,R} := \{z \in \mathbb{C} : r < z < R\}$, where $0 < r < R < \infty$.

Note that if $r/R = r'/R'$, then $A_{r,R}, A_{r',R'}$ are biholomorphic, via a scaling map; we will also use the notation $A_{r/R}$.

Hyperbolic annuli. The surface $A'_\lambda := \mathbb{H}/\langle z \mapsto \lambda z\rangle$, for any finite $\lambda > 0$, $\lambda \neq 1$, is also a topological annulus. Note that A'_λ is the same as $A'_{1/\lambda}$.

Proposition 2.1. *If $1 > \lambda_1 > \lambda_2 > 0$, then $A'_{\lambda_1}, A'_{\lambda_2}$ are not biholomorphically equivalent.*

Proof. If they were, the map would lift to a biholomorphic automorphism of universal covers $\mathbb{H} \to \mathbb{H}$. Since $Isom^+(\mathbb{H}) = \mathrm{Aut}(\mathbb{H})$, this map is an isometry, and it descends to an isometry $A'_{\lambda_1} \to A'_{\lambda_2}$. But the core geodesic of A'_λ has hyperbolic length $|\log \lambda|$, so the surfaces cannot be isometric, contradiction. ■

Proposition 2.2. $\mathbb{H}/\langle z \mapsto \lambda z\rangle \cong A_{\exp(-2\pi^2/\log\lambda)}$.

Proof. We first straighten out the semicircles with the map $f(z) = \log z$. The conjugate of $z \mapsto \lambda z$ by f is $z \mapsto z + \log \lambda$. We then apply the map $g(z) = \exp(\frac{2\pi i}{\log \lambda} z)$. ■

Proposition 2.3. *We have*

$$\mathbb{D}^* \cong \mathbb{H}/\langle z \mapsto z+1\rangle,$$
$$\mathbb{C}^* \cong \mathbb{C}/\langle z \mapsto z+1\rangle.$$

Proof. The (inverses of the) maps are both give by $z \mapsto \exp(2\pi i z)$. ■

Since the universal covers $\mathbb{H} \cong \mathbb{D}$ and $\mathbb{C}$ are not biholomorphic (by Liouville's theorem), $\mathbb{D}^*$ and $\mathbb{C}^*$ are not biholomorphic. And $\mathbb{C}^*$ is not biholomorphic to any A_r for the same reason. To see that $\mathbb{D}^*$ is not biholomorphic to any A_r we consider the hyperbolic geometry of $\mathbb{D}^* \cong \mathbb{H}/\langle z \mapsto z+1\rangle$; the core curve can be represented by an arbitrarily short hyperbolic curve.

Theorem 2.4. *Every Riemann surface that is homeomorphic to an annulus is biholomorphic to precisely one of $\mathbb{C}^*, \mathbb{D}^*$, or A_r, for some $0 < r < 1$.*

Proof. The "at most one" part of the theorem was proved in the paragraph above.

By the uniformization theorem, the universal cover is one of $\hat{\mathbb{C}}, \mathbb{C}, \mathbb{H}$. But the universal cover of an annulus is homeomorphic to the plane, which eliminates $\hat{\mathbb{C}}$.

An annulus covered by $\mathbb{C}$ is of the form $\mathbb{C}/\langle\gamma\rangle$, where γ is an automorphism of $\mathbb{C}$ without fixed points, i.e. a translation $z \mapsto z+a$. Any two such translations are conjugate by a scaling map, so these Riemann surfaces are all biholomorphic to $\mathbb{C}/\langle z \mapsto z+1\rangle \cong \mathbb{C}^*$.

An annulus covered by $\mathbb{H}$ is of the form $\mathbb{H}/\langle\gamma\rangle$, where γ is an automorphism of $\mathbb{H}$ without fixed points. Up to conjugation, there are two types of automorphism without fixed points, depending on whether the map fixes one or two points on the circle boundary of $\mathbb{H}$ (see Proposition 2.5 below.) The *hyperbolic* isometries are those conjugate to $z \mapsto \lambda z$, for some $\lambda > 0$. The *parabolic* isometries are those conjugate to $z \mapsto z \pm 1$. We get annuli biholomorphic to $\mathbb{H}/\langle z \mapsto \lambda z\rangle \cong A_{\exp(-2\pi^2/\log\lambda)}$, or $\mathbb{H}/\langle z \mapsto z+1\rangle \cong \mathbb{D}^*$, respectively. ■

2.1 Classification of isometries of $\mathbb{H}$

Proposition 2.5. *Every element of* $Isom^+(\mathbb{H})$ *is conjugate to precisely one of*

(i) $z \mapsto z \pm 1$,

(ii) $z \mapsto \lambda z$, *for some* $\lambda > 0$,

(iii) $f(z) = \frac{\cos\theta z + \sin\theta}{(-\sin\theta)z + \cos\theta}$, *for some* $\theta \in [0, \pi)$.

Proof. Part of the "at most one" part is seen by looking at the fixed points in $\bar{\mathbb{H}}$ (first one should note that all isometries extend to $\bar{\mathbb{H}}$). The number of fixed points is 1 (on the boundary), 2 (on the boundary), 1 (in the interior, assuming f is not the identity).

The other part can be proved using Jordan canonical form (with some care, since we are meant to conjugate by real matrices). ■

The trace (up to ± 1 factor) is often a helpful way of thinking about the classification: the cases correspond to $|tr| = 2, |tr| > 2, |tr| < 2$.

3 Teichmüller space and moduli space

Definition 3.1 (Moduli space)**.** Given a surface S, the moduli space $\mathcal{M}(S)$ is the set of all Riemann surfaces homeomorphic to S, up to biholomorphism.

A major goal in this course will be to understand this set, and the various structures it admits. If we want to talk about notions uniformly across different surfaces (such as a homotopy class of curves), we should use a more restrictive equivalence notion than biholomorphism.

Definition 3.2 (Teichmüller space)**.** Given an oriented topological surface S, the Teichmüller space is the set

$$\mathcal{T}(S) := \{(X, \phi) : X \text{ a Riemann surface, and } \phi : S \to X \text{ orient preserving homeo}\}/\sim,$$

where $(X, \phi) \sim (X', \phi')$ if there is a biholomorphic map $f : X \to X'$ with $f \circ \phi$ homotopic to ϕ'.

For $S = S_g$, the closed surface of genus g, we denote by $\mathcal{M}_g$ the set $\mathcal{M}(S)$; similarly for $\mathcal{T}$.

Example 3.3. For $S = A$ an annulus,

$$\mathcal{M}(A) = \{\mathbb{D}^*\} \cup \{A_r : 0 < r < 1\} \ \cup \{\mathbb{C}^*\},$$

$$\mathcal{T}(A) = \{(\mathbb{D}^*, id)\} \cup \{(\mathbb{D}^*, f)\} \cup \{A_r : 0 < r < 1\} \ \cup \{\mathbb{C}^*\},$$

where f is the unique (homotopy class of) homeomorphism of an annulus that is *not* homotopic to the identity id. On an annulus $A_{\lambda,1/\lambda}$ for some $0 < \lambda < 1$, such an f can be represented by the map $z \mapsto 1/z$. For the annuli other than $\mathbb{D}^*$, the markings by f and id are equivalent points in Teichmüller space, since f can be realized holomorphically on any such annulus; hence all markings for these are equivalent. For $\mathbb{D}^*$, the map f is not homotopic to a holomorphic automorphism (one can reduce the classification of $\mathrm{Aut}(\mathbb{D}^*)$ to that of $\mathrm{Aut}(\mathbb{D})$ using the Riemann removable singularity theorem).

Remark: Because of the subtleties discussed above, people don't usually bother too much with the Teichmüller space of the annulus.

Geometric structures on compact orientable Riemann surfaces.

1. The sphere. Such a Riemann surface is already simply connected, so by Uniformization Theorem, it must be $\hat{\mathbb{C}}$. The natural homogeneous geometry here is spherical. Both $\mathcal{M}_0$ and $\mathcal{T}_0$ are a single point. For the latter, use the fact that any orientation-preserving self homeomorphism of the sphere is homotopic to the identity.

2. Torus. By Uniformization Theorem, the universal cover $\hat{X}$ is biholomorphic to either $\hat{\mathbb{C}}, \mathbb{C}$ or $\mathbb{H}$. But we know the universal cover is homeomorphic to the plane, which rules out $\hat{\mathbb{C}}$. If universal cover were $\mathbb{H}$, X would admit a hyperbolic metric, but Gauss-Bonnet rules this out ($\chi(T) = 0$). Thus $\tilde{X} \cong \mathbb{C}$, and $X = \mathbb{C}/\Gamma$, where $\Gamma \subset \mathrm{Aut}(\mathbb{C}) = \{z \mapsto az+b : a, b \in \mathbb{C}, a \neq 0\}$. Not all elements of $\mathrm{Aut}(\mathbb{C})$ preserve the flat metric. But elements of Γ must be fixed-point free i.e. of the form $z \mapsto z + b$, and these do preserve the metric. So X admits a flat metric.

3. $S_g, g \geq 2$. As in previous case, $\tilde{X}$ must be $\mathbb{C}$ or $\mathbb{H}$. By same reasoning as above, if $\tilde{X} \cong \mathbb{C}$, then X admits a flat metric, but this is impossible by Gauss-Bonnet since $\chi(S_g) = 2 - 2g < 0$. Thus $\tilde{X} \cong \mathbb{H}$, and X admits a hyperbolic metric.

3.1 Teichmüller space of the torus

Definition 3.4. A *lattice* in $\mathbb{C}$ is a discrete subgroup of $\mathbb{C}$ isomorphic to $\mathbb{Z}^2$. A *marked positive lattice* is a lattice with a distinguished ordered pair of generators (γ_1, γ_2) that are positively oriented in $\mathbb{C}$, i.e. $\gamma_2/\gamma_1 \in \mathbb{H}$.

We define an equivalence relation $\sim$ on the space of marked positive lattices by $(\gamma_1, \gamma_2) \sim (\lambda\gamma_1, \lambda\gamma_2)$ for any $\lambda \in \mathbb{C}^*$.

Proposition 3.5. *There is a bijection*

$$\{\text{marked positive lattices}\}/\sim \quad \to \quad \mathcal{T}_1$$

given by $[(\gamma_1, \gamma_2)] \mapsto \mathbb{C}/\langle\gamma_1, \gamma_2\rangle$*, with the marking determined by the map on* $\pi(S) \cong \mathbb{Z}^2$ *given by* $(1, 0) \mapsto \gamma_1$ *and* $(0, 1) \mapsto \gamma_2$*.*

Proof. To define the inverse of the map, we start with X a marked Riemann surface structure on the torus and take some universal cover, which must be biholomorphic to $\mathbb{C}$. The deck group elements corresponding to $(1, 0), (0, 1) \in \mathbb{Z}^2 \cong \pi_1(S)$ act by translations γ_1, γ_2 on $\mathbb{C}$. Because the marking map is orientation preserving in the definition of Teichmüller space, we have that (γ_1, γ_2) is positively oriented. However, this description is not unique, since the picture can be conjugated by any holomorphic automorphism of $\mathbb{C}$. This can be taken to be $z \mapsto \lambda z$, $\lambda \in \mathbb{C}^*$, since translations commute with the deck group. Doing this gives the marked lattice $(\lambda\gamma_1, \lambda\gamma_2)$. Thus this gives a well-defined marked lattice up to rotation/scaling, and the resulting map is the inverse of the original. ■

Proposition 3.6. *There is a bijection*

$$\mathbb{H} \to \{\text{marked positive lattices}\}/\sim$$

given by $\tau \mapsto [(1, \tau)]$*.*

Proof. The inverse is given by $[(\gamma_1, \gamma_2)] \mapsto \gamma_2/\gamma_1$. ■

By combining the two bijections above, we get

Proposition 3.7. *There is a bijection*

$$\mathbb{H} \to \mathcal{T}_1$$

given by $\tau \mapsto \mathbb{C}/\langle 1, \tau\rangle$*, with the marking determined by the map on* $\pi(X) \cong \mathbb{Z}^2$ *given by* $(1, 0) \mapsto 1$ *and* $(0, 1) \mapsto \tau$*.*

The topology on $\mathcal{T}_1$ is defined so that this bijection is a homeomorphism.

The moduli space $\mathcal{M}_1$ is equal to the set of lattices modulo complex scaling. This can be understood as above, and then forgetting the marking, which means taking the quotient by $SL_2(\mathbb{Z})$. We get that $\mathcal{M}_1$ is in bijection with $\mathbb{H}/SL_2(\mathbb{Z})$, acting by Mobius transformations, which is the same as $\mathbb{H}/PSL_2(\mathbb{Z})$.

3.2 Mapping class group

Definition 3.8. Given a closed surface S, we define the mapping class group as

$$Mod(S) := \pi_0(\text{Homeo}^+(S)),$$

i.e. isotopy classes of homeomorphisms.

Fact 3.9. Let α, β be two simple closed curves on closed surface S. Then α, β are isotopic iff they are homotopic.

Fact 3.10. Let S be a closed surface, and $f, g \in \text{Homeo}(S)$. Then f, g are homotopic iff they are isotopic.

By the fact above, $Mod(S)$ is equal to $\text{Homeo}^+(S)$ modulo homotopy.

In dimension two (and three), any homeomorphism is isotopic to a diffeomorphism. Hence we also have

$$\pi_0(\text{Homeo}^+(S)) = \pi_0(\text{Diffeo}^+(S)).$$

Definition 3.11 (Dehn twist)**.** The Dehn twist of the annulus (which for concreteness we take to be $A_{1,2}$ is the map given in polar coordinates by

$$(r, \theta) \mapsto (r, \theta + 2\pi(r-1)).$$

Note that the Dehn twist extends to a map that is the identity on both boundary components.

On any surface S_g with a simple closed curve γ, we define the Dehn twist T_γ by taking a small annular neighborhood of γ and applying the above map on that region. On the rest of the surface T_γ is defined to be the identity.

For the torus S_1, we consider the action on the fundamental group

$$Mod(S_1) \to \text{Aut}(\pi_1(\mathbb{Z})) = \text{Aut}(\mathbb{Z}^2) \cong GL_2(\mathbb{Z}).$$

On general spaces/surfaces, one has to worry about base points, but since $\pi(S_1)$ is abelian, this is not an issue. We do always get an action on first homology $H_1(S;\mathbb{Z})$, and for the torus this actually coincides with the above, since in this case π_1 and H_1 are the same.

Since the maps are orientation preserving, induced maps on curves preserve algebraic intersection number, and so the image of the above lies in $SL_2(\mathbb{Z})$.

Proposition 3.12. *For the torus S_1, the map*

$$Mod(S_1) \to SL_2(\mathbb{Z})$$

is an isomorphism of groups.

Proof. For surjectivity, we consider linear automorphisms of the torus (by thinking of T as $\mathbb{R}^2/\mathbb{Z}^2$).

For injectivity, we use $K(\pi, 1)$ theory (i.e. Eilenberg-MacLane spaces): there is a bijection between endomorphisms of $\pi_1(S)$ and homotopy classes of maps $S \to S$. ∎

3.3 Action of mapping class group on Teichmüller space

For a marked surface (X, ϕ), we can pre-compose the marking with any homeomorphism $f : S \to S$ to get a new marked surface $(X, \phi \circ f)$. Because of the equivalence relation in terms of homotopy, the resulting point in Teichmüller space is well-defined, and depends only on the mapping class of f. Since we are precomposing, to get an action of Mod on Teichmüller space, we should take for any $f \in Mod(S)$,

$$[(X, \phi)] \mapsto [(X, \phi \circ f^{-1})].$$

The action is by changing markings; it follows that the quotient "forgets" the marking, hence

$$\mathcal{M}(S) = \mathcal{T}(S)/Mod(S).$$

Proposition 3.13. *With $\mathcal{T}(S_1)$ identified with the torus, the action of $Mod(S) \cong SL_2(\mathbb{Z})$ is given by*

$$\tau \mapsto \frac{a\tau - b}{-c\tau + d}$$

(It differs from the standard projective linear action on $\mathbb{H}$ because of (i) the inverse in the definition of the action, and (ii) the choice to normalize the image of $(1, 0)$, rather than $(0, 1)$, to be the resulting point in $\mathbb{H}$).

Note that the action is not faithful: the matrix $-I$ acts trivially. This fact is equivalent to the fact that every complex automorphism admits a holomorphic automorphism in the mapping class of the *hyperelliptic involution* in $Mod(S_1)$ (the map corresponding to the $-I$ matrix).

3.4 Fundamental domain for $\mathcal{M}_1$

Proposition 3.14. *A fundamental domain for the action of $SL_2(\mathbb{Z})$ on $\mathbb{H}$ is given by the region*

$$\{z \in \mathbb{H} : -1/2 \leq \mathrm{Re}(z) \leq 1/2, \quad |z| \geq 1\}.$$

Proof. Given τ, we begin by assuming it has maximal imaginary part among all elements in its $SL_2(\mathbb{Z})$ orbit (the maximum is attained, which can be seen using the formula $\mathrm{Im}\left(\frac{az+b}{cz+d}\right) = \frac{\mathrm{Im}(z)}{|cz+d|^2}$, found by multiplying denominator by it's conjugate).

Now the group is generated by the matrices corresponding to $z \mapsto z + 1$ and $z \mapsto -1/z$ (as can be seen by doing the Euclidean algorithm in the first column, to make the lower-left entry zero). By applying the first, we can bring τ to the strip $\{-1/2 \leq \mathrm{Re}(z) \leq 1/2\}$. If $|\tau| \leq 1$, then by applying $z \mapsto -1/z$, we get a point with larger imaginary part than the original, contradiction. Hence we can move any point in $\mathbb{H}$ into the domain.

By examining what $z \mapsto z + 1$ and $z \mapsto -1/z$ do, we see that our set is in fact a fundamental domain. ■

We should also study the identifications on the boundary. The left and right vertical sides get glued together, while the left and right halfs of the bottom arc are glued. The resulting complex orbifold is known as the modular curve. (As a topological orbifold, it is the $(2, 3, \infty)$ orbifold).

Fixed points of action: not free, but properly discontinuous.

Proposition 3.15. *The action of $PSL_2(\mathbb{Z})$ on $\mathbb{H}$ is* properly discontinuous: *for each $K \subset \mathbb{H}$ compact, the set*

$$\{g \in PSL_2(\mathbb{Z}) : gK \cap K \neq \emptyset\}$$

is finite.

Proof. This can be proved using the fundamental domain for the action constructed above. ■

We will see that in higher genus, though there is no such explicit fundamental domain, the action of the mapping class group on Teichmüller space will still be properly discontinuous.

Fixed points. By looking at the definition of Teichmüller space and the action of the mapping class group, we find that a fixed point of some $f \in Mod(S_g)$ corresponds to a Riemann surface with a biholomorphism in the homotopy class of (conjugate) of f.

The fixed points $i, e^{2\pi i/3}$ in $\mathbb{H}$ correspond to special tori. The point i corresponds to the square torus. This admits an automorphism corresponding to rotation by $\pi/2$.

The point $e^{2\pi i/3}$ corresponds to hexagonal torus (a regular hexagon with opposite sides glued together; this can be cut up into the parallelogram given by $1, e^{2\pi i/3}$). The lattice is easily seen to have $\pi/3$ rotational symmetry, and this symmetry is also evident from the hexagon description.

3.5 Fenchel-Nielsen coordinates

A *pair of pants* is just a sphere with 3 boundary components, the cuffs.

Proposition 3.16. *Every surface of genus $g \geq 2$ can be cut into $2g - 2$ pairs of pants.*

Proof. Pick a simple closed curve on the surface that is not homotopic to a point. In the complement find another such curve, and continue doing this, until one cannot proceed. The complementary regions will all be pairs of pants. The Euler characteristic of a pair of pants is -1, so by additivity of Euler characteristic, there must $\chi(S_g)/-1 = 2g - 2$ of them. ■

Proposition 3.17. *Given $\ell_1, \ell_2, \ell_3 \in \mathbb{R}_{>0}$, there is a unique hyperbolic structure on a pair of pants such that the cuffs are geodesics with lengths ℓ_1, ℓ_2, ℓ_3.*

Proof. For each pair of cuffs, we find the shortest arc from one to the other, the *seams* of the pants. This will be a geodesic arc meeting each of the cuffs at right angles (otherwise, there would be a shortcut). This decomposes the pants into two right-angled hyperbolic hexagons. By considering the seams, which are each a side of both hexagons, the proposition below implies the two hexagons are isometric. Hence each of the cuffs of the original pants is bisected by the seams. And applying the proposition below again, we see that there is exactly one choice for the hexagon, since alternating sides must have lengths $\ell_1/2, \ell_2/2, \ell_3/2$. ∎

Proposition 3.18. *Given $\ell_1, \ell_2, \ell_3 \in \mathbb{R}_{>0}$, there is a unique hyperbolic right-angled hexagon (all sides geodesic) such that* alternating *sides have lengths ℓ_1, ℓ_2, ℓ_3.*

For a soft proof, see Thurston-Levy book Fig 4.15; it can also be proved using hyperbolic trigonometry.

Length functions. Give a closed curve γ on our reference surface S_g, we get a length function

$$\ell_\gamma : \mathcal{T}_g \to \mathbb{R}_{\geq 0}$$

that takes X to the hyperbolic length on X of (the geodesic representative) of γ on X. If γ is an essential curve (not homotopic to a point), then the length is positive.

Defining the coordinates. We will define a map

$$FN : \mathcal{T}_g \to \mathbb{R}^{3g-3}_{>0} \times \mathbb{R}^{3g-3}.$$

Fix a pants decomposition $\mathcal{P}$ on our reference surface S_g. Let $c_1, \ldots, c_{3g-3}$ be the cuff curves. The first three components of the map are the *length* parameters, given by $\ell_{c_1}(X), \ldots, \ell_{c_{3g-3}}(X)$.

The rest of the components of the map are *twisting* parameters. Intuitively, τ_i measures the amount of twisting that is done when gluing the pairs of pants on either side of c_i to one another. Because we are working with Teichmüller space, a full turn results in a different marked surface, and so it turns out we will get a well-defined real number (rather than say an angle) as the twist parameter.

To define the twist formally, for each pair of pants in $\mathcal{P}$, we take disjoint (oriented) arcs connecting each pair of cuffs. We then join choose some pairing of arcs coming into each cuff from the pants on each either side. This gives a collection $T = \{t_1, \ldots, t_k\}$ of (oriented) simple closed curves on the surface. This allows us to define the twist parameter τ_i as the twisting of T around γ_i relative to the seams of the pants on X; this is a real number, and it is normalized such that the full Dehn twist about γ_i will increase the τ_i by $\ell_{\gamma_i}(X)$. (Given an isotopy class of arc on a pants with endpoints on cuffs, the twist of the arc around each cuff relative to each seam is well-defined up to isotopy of the arc fixing the basepoints. To define the twist about a cuff in the above, we find the two arcs that form curves in T, and measure their twists around the cuff relative to the corresponding seam, and then take the difference between these two numbers).

Dimension counts. The real dimension $6g-6$ of Teichmüller space can also be informally derived in other ways. For instance, for $g=2$, we can make a surface by gluing together opposite sides of a regular octagon with all 45 degree angles. So:

$+16$ dimension worth of choices for the vertices of the octagon.

-3 dimensions of isometry group of $\mathbb{H}$

-4 conditions on opposite sides having equal length

-1 condition on angle at vertex of glued surface being 360.

-2 dimensions worth of choice of the point that becomes vertex of octagon.

This gives $16-3-4-1-2=6$, which agrees with $6(2)-6$.

3.6 Character variety

Fundamental group of surfaces. S_g can be realized as a $4g$-gon with sides glued according to the pattern

$$\alpha_1\beta_1\alpha_1^{-1}\beta_1^{-1}\cdots\alpha_g\beta_g\alpha_g^{-1}\beta_g^{-1}.$$

This leads to the following presentation of the fundamental group:

$$\pi_1(S_g)=\langle\alpha_1,\beta_1,\ldots,\alpha_g,\beta_g \mid \alpha_1\beta_1\alpha_1^{-1}\beta_1^{-1}\cdots\alpha_g\beta_g\alpha_g^{-1}\beta_g^{-1}\rangle.$$

One should think of it as somewhere between a free abelian group on $2g$ generators, and a free group on $2g$ generators. In the torus case $g=1$, it coincides with the free abelian group.

Let $DF(\pi_1(S_g), PSL_2(\mathbb{R}))$ denote the set of faithful homomorphisms $\pi_1(S_g)\to PSL_2(\mathbb{R})$ with discrete image. These are exactly the "covering space actions" of $\pi_1(S_g)$ on $\mathbb{H}$ by isometries, i.e. those corresponding to compact hyperbolic surfaces.

Proposition 3.19. *There is a bijection between $\mathcal{T}_g$ and the character variety*

$$DF(\pi_1(S_g), PSL_2(\mathbb{R}))/PGL_2(\mathbb{R}),$$

where the quotient denotes the action by conjugation.

Proof. Starting with a point in Teichmüller space, choose a universal cover. The fundamental group of X acts by isometries on this, giving the representation. It is discrete and faithful.

For the other direction, take quotient. (To get actual marking, need to use that any (outer) automorphism of $\pi_1(S_g)$ is induced by a homotopy equivalence $S\to S$, which in turn is induced by a homeomorphism). ■

Algebraic topology on $\mathcal{T}_g$. As a space of continuous functions, $Hom(\pi_1(S_g), PSL_2(\mathbb{R}))$ (we give $\pi_1(S_g)$ the discrete topology) has a compact-open topology. As a quotient of a subset of this, the character variety also inherits a topology. We can also see this concretely, by embedding $Hom(\pi_1(S_g), PSL_2(\mathbb{R}))$ into $(PSL_2(\mathbb{R}))^{2g}$ via images of generators (and then taking quotient of subspace).

Dimension count. We can also get a guess for dim $\mathcal{T}_g$ using this perspective:

$+(3)(2g)$ dimension choices of images of the $2g$ generators in $PSL_2\mathbb{R}$.

-3 condition that product of commutators is I

-3 dimensions for conjugating by $PGL_2(\mathbb{C})$.

This gives $6g-6$, agreeing with what we get from Fenchel-Nielsen coordinates.

3.7 $9g-9$ theorem

Proposition 3.20. *There is a set of* $9g-9$ *simple closed curves* $\alpha_1, \ldots, \alpha_{9g-9}$ *whose lengths determine a point in Teichmüller space.*

Proof idea. To choose the curves, begin with the $3g-3$ cuff curves of a pants decomposition. For each such cuff α, add some γ that crosses α at least once and does not cross any other cuffs. Also add each $\gamma = T_\alpha(\gamma)$.

The length Fenchel-Nielsen coordinates with respect to this pants decomposition can then be immediately recovered. We claim the twist about α can be recovered from the lengths of γ, γ. It is a fact that if we fix all the cuff lengths, then the lengths $\ell_t(\gamma), \ell_t(\gamma')$ is a strictly convex function of the twist t about α. An individual convex function need not be injective. However the lengths of γ and γ' are just shifts of each other by $\ell(\alpha)$, since they are related by a full Dehn twist. The graph of a convex function cannot have two horizontal chords of equal length, hence $t \mapsto (\ell_t(\gamma), \ell_t(\gamma'))$ is injective.

■

Proposition 3.21. *Any Fenchel-Nielsen coordinate map*

$$FN : \mathcal{T}_g \to \mathbb{R}^{3g-3}_{>0} \times \mathbb{R}^{3g-3}$$

is a homeomorphism with respect to algebraic topology on domain and standard Euclidean topology on target.

Proof. Since we've already shown the map is a bijection, by invariance of domain it suffices to show that it's continuous. The length parameters are continuous functions, since the length of a closed curve is a continuous function of the squared trace of the image of an element of π_1. Any twist can be expressed as continuous functions of other lengths, as in proof of $9g-9$ theorem, so the twist parameters are also continuous. ■

Aside: genus two Riemann surfaces are hyperelliptic. A hyperelliptic involution of a surface is a mapping class represented by an order two homeomorphism with $2g+2$ fixed points (so the quotient map is a degree two cover to the sphere branched over $2g+2$ points, by Riemann-Hurwitz formula). It can be represented by linearly arranging the surface, "skewering", and rotating by 180 degrees. In genus two, this map fixes every isotopy class of simple closed curve (though it changes orientation of some curves). By the $9g-9$ theorem, this implies that this mapping class acts trivially on Teichmüller space. So the mapping class group action in genus two is not faithful.

4 Teichmüller existence and uniqueness theorems

4.1 Quasiconformal maps

Our goal in this section is to define a notion of distance between Riemann surfaces. There are several sensible ways of doing this. From the perspective of complex analysis, the most natural way leads to study of quasi-conformal maps and the Teichmüller metric.

Complex derivatives and dilatation. Given a smooth map $f : U \to \mathbb{R}^2$, where $U \subset \mathbb{R}^2$ is an open set containing 0, we wish to measure the failure to be conformal at 0. We write $f(x,y) = \begin{pmatrix} u(x,y) \\ v(x,y) \end{pmatrix}$, and $Df = \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix}$. The matrix is conformal iff it is a composition of rotation and scaling iff $u_x = v_y$ and $v_x = -u_y$. In terms of differentials, and identifying $\mathbb{R}^2$ with $\mathbb{C}$ we can write

$$df = (u_x + iv_x)dx + (u_y + iv_y)dy = f_x dx + f_y dy.$$

Instead of the 1-forms dx, dy, we wish to express in terms of $dz = dx + idy$ and $d\bar{z} = dx - idy$. Solving for dx, dy in terms of these and substituting in above gives

$$df = \frac{1}{2}(f_x - if_y)dz + \frac{1}{2}(f_x + if_y)d\bar{z}.$$

Motivated by this, we set $f_z = \frac{1}{2}(f_x - if_y)$, $f_{\bar{z}} = \frac{1}{2}(f_x + if_y)$. The Cauchy-Riemann equations are equivalent to $f_{\bar{z}} = 0$.

We define the *complex dilatation* of f (at some point p) to be

$$\mu_f(p) := \frac{f_{\bar{z}}}{f_z}.$$

(This is better than using just $f_{\bar{z}}$, which isn't invariant under scaling).

By relating the determinant of Df to $f_z, f_{\bar{z}}$, one can check that f is orientation preserving at 0 iff $\mu_f < 1$.

We will generally work with a related quantity that is more closely connected to ratios of lengths. For an orientation preserving homeomorphism $f : U \to V$ between domains in the plane, we define the *dilatation* at a point p

$$K_f(p) := \frac{1 + |\mu_f(p)|}{1 - |\mu_f(p)|}.$$

Geometric interpretation. At a point, (the derivative) of such a smooth map takes a circle centered at 0 in the tangent space to an ellipse. The dilatation is the ratio of the major to minor axes of this ellipse. (In other words, it is the ratio of singular values of the matrix).

Dilatation of a smooth map. We define

$$K_f := \sup_{p \in U} |K_f(p)|.$$

The following are easy to verify:

- $K_{f \circ g} \leq K_f K_g$ (f, g defined on suitable domains)
- $K_{f^{-1}} = K_f$.
- $K_{f \circ g} = K_f = K_{g \circ f}$ if g is conformal.

The last property means that K_f is invariant under conformal change of variable.

We say f is K-*quasiconformal* if $K_f \leq K$. We say f is *quasiconformal* if it is K-quasiconformal for some K.

Quasiconformal maps of Riemann surfaces. We can extend all these notions to a map $f : X \to Y$ between Riemann surfaces. One should not insist that quasiconformal maps are smooth everywhere. We will use the condition that the map is smooth except at finitely many points (and then K_f is defined as a sup of $K_f(p)$ over only the smooth points). The actual general definition of quasiconformal maps insists on much less regularity, but that would take us into lots of analysis.

Proposition 4.1. *Such an $f : X \to Y$ is conformal iff it is* 1*-qc.*

Proof. We see that the map is conformal away from the points where f is not smooth. Then since f is a homeomorphism, by removable singularity theorem, the map is actually conformal at these points as well. ■

Goal. Given such an f, find the map in its homotopy class that minimizes the dilatation.

Teichmüller metric. We let

$$d_{Teich}(X,Y) = \log \inf_f K_h,$$

where the infimum is over all f in the homotopy class of homeomorphisms $X \to Y$ coming from the markings of X, Y. We can see already that f is a pseudometric, but it is not yet clear that distances between different points are always positive.

4.2 Grotzsch inequality

We investigate the goal above first in a situation without topology. This is the key result that will allow us to understand optimal conformal maps.

Theorem 4.2 (Grotzsch)**.** *Let $F : [0,a] \times [0,1] \to [0, Ka] \times [0,1]$, $K \geq 1$, be an orientation preserving homeomorphism, smooth away from finitely many points, and that takes horizontal (resp vertical) sides to horizontal (resp vertical) sides. Then*

$$K_F \geq K,$$

with equality iff K is affine.

Proof. The idea is that F must stretch in the horizontal direction, but if it also stretched a lot in the vertical direction, then it would expand area too much. Thus we can bound the ratio of maximal to minimal stretch from below.

From the Fundamental Theorem of Calculus, we get that, for each fixed y,

$$Ka \leq \int_0^a \|F_x(x,y)\| dx.$$

Integrating this over y yields

$$Ka = \int_0^y Ka \leq \int \|F_x(x,y)\| dA.$$

Now $F_x(x,y)$ is the image of a unit tangent vector under the derivative mapping, so it's magnitude is less than the maximum stretch M of F at (x,y). If the minimum stretch is m, note that $K_F(x,y) = M/m$, while the Jacobian $J_F(x,y)$ equals mM. Hence $\|F_x(x,y)\| \leq M = \sqrt{K_F(x,y)J_F(x,y)}$. Using this and then Cauchy-Schwarz yields

$$\begin{aligned} Ka &\leq \int \sqrt{K_F(x,y)J_F(x,y)} dA \\ &\leq \left(\int K_F(x,y) dA \int J_F(x,y) dA \right)^{1/2} \\ &\leq \left(a(K_F)(Ka) \right)^{1/2}, \end{aligned}$$

whence $K_F \geq K$.

If equality holds here, it must hold in all the inequalities above. The first implies the map takes horizontal lines to horizontal lines. Equality in the inequality for M implies the direction of maximal stretch is the horizontal direction. Equality in Cauchy-Schwarz implies $K_F(x,y)/J_F(x,y) = 1/m^2$ is constant, while the last inequality implies $K_F(x,y)$ is constant, hence M is constant. Since direction of minimal stretch is orthogonal to max stretch direction, it must be vertical. Hence the derivative of F is constant, i.e. it is affine. ■

4.3 Quadratic differentials

We would like to use Grotzsch's theorem locally on a Riemann surface to understand quasiconformal maps, and in particular the foliations along which they stretch maximally (or minimally). However, such foliations (and their tangent line fields) must have singularities; this is related to the Poincare-Hopf formula for vector fields: $\chi(M) = \sum_v \operatorname{index}(v)$.

A quadratic differential is a section of the square of the holomorphic cotangent bundle of the Riemann surface. Concretely it is a tensor that locally has the form $q(z) = f(z)dz^2$, where f is a holomorphic function. That is, under a coordinate change $z = \phi(w)$, we have

$$q = f(\phi(w))d(\phi(w))^2 = f(\phi(w))\phi'(w)^2 dw^2.$$

Construction of quadratic differentials: Start with a (collection of) polygon in the plane, each side of which has a partner side that is parallel and of equal length. Glue together partner sides by translation, or translation composed with 180 degree rotation. These preserve the form dz^2 on $\mathbb{C}$, so descend to give quadratic differentials on the surface. The term *half-translation surface* is often used to describe this geometric structure (the "half" is meant to suggest rotation by 180 degrees).

To find Riemann surface charts for the quotient surface, one uses Euclidean translation maps, away from the corners of the polygons. At the corners, one uses maps of the form $z \mapsto z^{k/2}$.

Examples: regular octagon with opposite sides identified, swiss cross. The square pillowcase corresponds to a *meromorphic* quadratic differential (at some points it looks like $(1/z)dz^2$).

Zeros and cone points. A quadratic differential on a higher genus Riemann surface must have zeros; the sum of the zeros with multiplicities is $4g-4$. The analogous fact about half-translation surfaces is that the sum of the excess cone angle (angle at a point minus 2π) equals $\pi(4g-4)$. One sees that pulling back the form dz^2 to the coordinate charts gives a quadratic differential with a zero at the point coming from the cone point. For instance, for a 3π cone angle point

$$[d(z^{3/2})]^2 = [(3/2)z^{1/2}dz]^2 = (9/4)z\,dz^2,$$

which has a zero (of order 1).

Measured foliations. A quadratic differential gives a foliation of the surface in every direction (in particular, in horizontal and vertical directions), except that it can have singularities at the zeros of the differential.

From the polygon perspective, the foliations come from e.g. the horizontal foliation of the Euclidean plane; it descends to a (singular) foliation of the surface since translations and 180 degree rotation preserve the foliation.

From the complex analytic perspective, the horizontal foliation is the one tangent to the vectors which are assigned a positive real value by q (one can think of a quadratic differential as giving a $\mathbb{C}$-valued function to very tangent space to the surface, which is the square a of a linear function, and which varies holomorphically over the base point).

In fact, each foliation comes with a *tranverse measure*, i.e. a measure μ on arcs transverse α to the foliation, which is invariant under homotopies preserving the leaves. For the horizontal foliation, this is given by

$$\mu(\alpha) = \int_\alpha |\operatorname{Im}(\sqrt{q})|.$$

Note that the expression $\sqrt{q}$ does not make sense globally, but locally (away from zeros of the differential), one can choose a branch of the square root to get a 1-form. For dz^2, the transverse measure for the horizontal foliation is just vertical Euclidean distance (the form integrated is just $|dy|$).

Coordinate charts from quadratic differential. The additional information of the quadratic differential q determines some particularly nice "natural coordinates" charts for the Riemann surface X. Fix some point $p \in X$. In a small neighborhood of p, we define a function

$$g(z) = \int_p^z \sqrt{q},$$

choosing some branch of the square root. This is a holomorphic function (since all operations involved are "holomorphic"), and its derivative at p is non-zero, so it gives a local homeomorphism near p.

We claim that the pullback of dz^2 under g is q. Indeed, by the fundamental theorem of calculus

$$g^*(dz^2) = d(g(z))^2 = (\sqrt{q})^2 = q.$$

So these are the coordinates in which the differential "looks like" dz^2, i.e. looks like the Euclidean plane.

Dimension of quadratic differentials.

Proposition 4.3. *Let $QD(X)$ be the complex vector space of holomorphic quadratic differentials on X. Then*

$$\dim_{\mathbb{C}} QD(X) = 3g - 3.$$

This follows from the Riemann-Roch Theorem.

Since there is a canonically defined vector space associated to each point in Teichmüller space, we get a rank $3g-3$ vector bundle over Teichmüller space. One can also compute the dimension of the total space by counting independent parameters in the polygon construction.

4.4 Teichmüller uniqueness theorem

Teichmüller maps

Definition 4.4. Given X Riemann surface, q a quadratic differential, and $K \geq 1$, we define a new Riemann surface Y equipped with natural charts obtained from those for (X, q) by composing with the map $f(x+iy) = \sqrt{K}x + i\frac{1}{\sqrt{K}}y$; these give rise to a quadratic differential q' on X'. Furthermore, we get a map $h : (X, q) \to (Y, q')$ called the *Teichmüller* map associated to (X, q, K).

Example: We can take a polygon representation and apply the map f to it.

Theorem 4.5 (Teichmüller uniqueness)**.** *Let $h : X \to Y$ be a Teichmüller map with data (X, q, K). Then h is the unique map in its homotopy class with minimal dilatation.*

That is, any other quasiconformal map $f : X \to Y$ homotopic to f satisfies $K_f > K$.

This is an analog of Grotzsch's theorem, with a rectangle replaced by a Riemann surface. We will need a few preliminary lemmas before the proof. We measure norms of derivatives with respect to the flat coordinates induced by q on X and q' on Y.

Lemma 4.6. *We have*

$$\int_X \|f_x\| dA_q \geq \sqrt{K}\,\mathrm{Area}(q).$$

Proof. Define $\delta : X \times \mathbb{R}_{\geq 0} \to R_{\geq 0}$ by

$$\delta(p, L) = \int_{\alpha_{p,L}} \|f_x\| dx,$$

where $\alpha_{p,L}$ is the horizontal arc of length $2L$ centered at p. We have

$$\delta(p, L) = \ell_{q'}(f(\alpha_{p,L})).$$

Note that $\delta(p, L)$ is not defined for large L when p is on a critical trajectory (one heading to a cone point). However this is a measure 0 set, so will not affect the following.

Integrating over p and using Fubini gives

$$\int_X \delta(p, L) dA = 2L \int_X \|f_x\| dA.$$

On the other hand, by the Lemma below (applied with map $f \circ h^{-1}$ and arc $h(\alpha)$, which has length $2L\sqrt{K}$):

$$\begin{aligned}\int_X \delta(p, L) dA &= \int_X \ell_{q'}(f(\alpha_{p,L})) dA \\ &\geq \int_X (2L\sqrt{K} - C) dA \\ &= (2L\sqrt{K} - C) \operatorname{Area}(q).\end{aligned}$$

Combining with the previous equality gives

$$\int_X \|f_x\| dA \geq (\sqrt{K} - \frac{C}{2L}) \operatorname{Area}(q),$$

and then taking $L \to \infty$ yields the desired result.

■

Lemma 4.7. *Let (X, q) be a Riemann surface with quadratic differential. Let $f : X \to X$ be a homeomorphism that's isotopic to the identity. There exists some constant C such that for any arc $\gamma : [0, 1] \to X$ in a leaf of a horizontal foliation*

$$\ell_q(f(\alpha)) \geq \ell_q(\alpha) - C.$$

Proof. The homotopy moves the endpoints of arcs by a bounded amount; take C to be twice this amount. Then the result follows from triangle inequality, since α achieves the minimum length in its homotopy class rel endpoints (this would fail if q had simple poles, and homotopies were considered on the unpunctured surface; for instance cylinders on a rectangular pillowcase are null-homotopic).

■

Proof of Theorem 4.5 (Teichmüller uniqueness). With the lemmas we've developed, we can now follow the proof of Grotzsch's inequality.

By Lemma 4.6,

$$\sqrt{K} \operatorname{Area}(q) \leq \int_X \|f_x\| dA_q$$

Now $f_x(x, y)$ is the image of a unit tangent vector under the derivative mapping, so its magnitude is less than the maximum stretch M of f at (x, y). If the minimum stretch is m, note that $K_f(x, y) = M/m$, while the Jacobian $J_f(x, y)$ equals mM. Hence $\|f_x(x, y)\| \leq M = \sqrt{K_f(x, y) J_f(x, y)}$. Using this and then Cauchy-Schwarz yields

$$\begin{aligned}
\sqrt{K}\,\mathrm{Area}(q) &\le \int_X \sqrt{K_f(x,y)J_f(x,y)}dA_q \\
&\le \left(\int_X K_f(x,y)dA_q \int_X J_f(x,y)dA_q\right)^{1/2} \\
&\le \left((\mathrm{Area}(q)\cdot K_f)\cdot(\mathrm{Area}(q'))\right)^{1/2} \\
&\le \left((\mathrm{Area}(q)\cdot K_f)\cdot(\mathrm{Area}(q))\right)^{1/2},
\end{aligned}$$

whence $K_f \ge K$.

The analysis of the equality case is exactly the same as in Grotzsch's inequality. ■

4.5 Teichmüller existence theorem

Theorem 4.8 (Teichmüller existence)**.** *Let $X, Y \in \mathcal{T}_g$. There exists a Teichmüller map $h : X \to Y$.*

Proof. The *area* $\|q\|$ of q is given by $\int_X |q|$. We let $QD(X)_1$ denote the space of quadratic differentials of area less than 1. We define a map $\Phi : QD(X)_1 \to \mathcal{T}_g$ that takes q to the image of the Teichmüller map with data $(X, q, \frac{1+\|q\|}{1-\|q\|})$. To prove the theorem, it suffices to show that Φ is surjective. To do this, we will use invariance of domain, below.

Continuity of Φ is reasonable, since if the polygon picture associated to q changes slightly, the one associated to q' only changes slightly. (This can also be justified using the Measurable Riemann Mapping Theorem and Beltrami differentials, ideas that lead to the Bers embedding.)

Injectivity of Φ follows from Teichmüller uniqueness theorem: if $\Phi(q_1) = \Phi(q_2) = Y$, then q_1, q_2 give two different homotopic Teichmüller maps, which minimize dilatation, but there is a unique map in the homotopy class minimizing dilatation.

To see properness of Φ, suppose $K \subset \mathcal{T}_g$ is compact. Note that $Y' \mapsto d_{Teich}(X, Y')$ is a continuous function of Y', fixing X (though recall that we don't know that $d_{Teich}(X, Y')$ is a metric yet). This is because if Y', Y'' are close, we can represent them with close by hyperbolic fundamental domains, and in particular there is a smooth map from one to the other with derivative close to the identity, and hence dilatation close to 1. It follows that $Y' \mapsto d_{Teich}(X, Y')$ is bounded, say by R, on K. If $\Phi^{-1}(K)$ were not compact, there would exist $q_1, q_2, \ldots$ with $\|q_i\| \to 1$ and $d_{Teich}(X, \Phi(q_i)) \le R$. But by definition there is a Teichmüller map from X to $\Phi(q_i)$ with dilatation $\frac{1+\|q_i\|}{1-\|q_i\|}$. By the Teichmüller uniqueness theorem, this is the map of minimal dilatation in its homotopy class, so $d_{Teich}(X, \Phi(q_i)) \ge \frac{1+\|q_i\|}{1-\|q_i\|} \to \infty$ as $i \to \infty$ contradiction.

Finally, we know that both $QD_1(X)$ and $\mathcal{T}_g$ are homeomorphic to $\mathbb{R}^{6g-6}$ (Proposition 4.3 and Proposition 3.21, respectively). Since Φ is continuous, injective, and proper between these, invariance of domain implies it is a homeomorphism, and in particular surjective. ■

Theorem 4.9 (Invariance of domain)**.** *A continuous, proper, injective map* $\mathbb{R}^n \to \mathbb{R}^n$ *is a homeomorphism.*

4.6 Teichmüller metric

Proposition 4.10. *The function* $d_{Teich}(X,Y)$ *defines a complete metric on Teichmüller space.*

Proof. Symmetry and triangle inequality follow from basic properties of quasiconformal maps. Positive definiteness follows from Teichmüller's existence theorem: if $X \neq Y$, there is a Teichmüller map $h : X \to Y$ with dilatation $K > 1$, and by Teichmüller's uniqueness theorem, this minimizes dilatation in the homotopy class, so $d_{Teich}(X,Y) = \log K > 0$.

For completeness, it suffices to show that any closed metric ball $B(X,R)$ is compact. The Teichmüller existence and uniqueness theorems imply that the map $\Phi : QD_1(X) \to \mathcal{T}_g$ restricts, for each $0 < s < 1$ to a homeomorphism $QD_s(X) \to B(X, \log \frac{1+s}{1-s})$. Note that $QD_s(X)$ is compact since it is the norm ball in a finite dimensional normed vector space. Thus $B(X, \log \frac{1+s}{1-s})$ also is compact, and by varying s we get all balls. ■

A *geodesic* in a metric space is an isometric embedding of $\mathbb{R}$ into the space. (When we work with spaces with more interesting topology, we will often take a geodesic to be a *local* isometric embedding). A *geodesic segment* is an isometric embedding of some interval $[a,b]$.

Proposition 4.11. *Every geodesic segment in* $(\mathcal{T}_g, d_{Teich})$ *is contained in a Teichmüller geodesic. Furthermore, between any two points* $X, Y \in \mathcal{T}_g$*, there is a unique geodesic segment.*

5 Nielsen-Thurston classification

5.1 Genus 1

Recall that the mapping class group of the torus is identified with $SL_2(\mathbb{Z})$, which acts on the upper-half plane $\mathbb{H}$ via Mobius transformations. We've already classified elements of $Isom^+(\mathbb{H})$ up to conjugacy into three classes. We'll think about this classification, restricted to those isometries coming from $f \in SL_2(\mathbb{Z})$, in terms extension of the action of Mod to the boundary $\partial\mathbb{H}$.

1. *Elliptic* elements are those that fix a point in $\mathbb{H}$, and are thus rotation. Since $SL_2(\mathbb{Z})$ is discrete, the rotation is *finite order*. The order must be $2, 3, 4, 6$.

2. *Parabolic* elements are those that fix exactly one point in $\partial\mathbb{H}$. This means that there is a unique real eigenvector. Since eigenvalues come in conjugate pairs, the other eigenvalue is also real. Since their product is 1, the eigenvalue has to be ± 1, and it occurs with multiplicity 2. The eigenvector is a solution to a linear system with integer coefficients, hence it can be taken to be rational. This means that f fixes an (unoriented) simple closed curve; we call f *reducible*. Any such f is either a power of a Dehn twist, or the hyperelliptic involution composed with a power of a Dehn twist.

3. *Hyperbolic* elements are those that fix a pair of points in $\partial\mathbb{H}$. Each corresponding real eigenvector gives a pair of invariant foliations on the torus. The eigenvalues are $\lambda, 1/\lambda$. We can apply a real matrix A that takes the pair of eigenvectors to the standard basis vectors. This conjugates f to $diag(\lambda, 1/\lambda)$. This map is a Teichmüller mapping on the torus T' with lattice generated by columns of A. The initial quadratic differential just comes from dz^2. The image surface is T', but with its marking composed with the mapping class.

Definition 5.1. An *Anosov* homeomorphism f of the torus to itself is one for which there exists a pair of (nonsingular) measured foliations $(\mathcal{F}_s, \mu_s)$, $(\mathcal{F}_u, \mu_u)$ such that

1. $\mathcal{F}_s, \mathcal{F}_u$ are transverse

2. $f(\mathcal{F}_u, \mu_u) = (\mathcal{F}_u, \lambda\mu_u)$ and $f(\mathcal{F}_s, \mu_s) = (\mathcal{F}_s, \lambda^{-1}\mu_s)$.

A mapping class is called Anosov if it has an Anosov representative. The number λ is called the *stretch factor*.

In the classification above, for a hyperbolic element, pulling back the vertical and horizontal foliations from T' to T gives such an Anosov package. Hence the above can be summarized as:

Proposition 5.2. *Every non-trivial mapping class on the torus is either periodic, reducible (fixes an isotopy class of simple closed curve), or Anosov.*

Translation length. Instead of thinking of the action on the boundary, we can instead look at the translation length:

$$\tau(f) := \inf_{z \in \mathbb{H}} d_{\mathbb{H}}(z, f(z)).$$

We then have

1. Elliptic corresponds to $\tau(f) = 0$, and infimum is realized

2. Parabolic corresponds to $\tau(f) = 0$, and infimum not realized.

3. Hyperbolic corresponds to $\tau(f) > 0$ (and realized).

Our proof in the higher genus case will make use of translation distance rather than a boundary for Teichmüller space.

5.2 Higher genus

Finite order. Examples of finite order mapping classes, can be obtained by realizing the surface in three-dimensional space, and rotating. Or realize the surface as regular $4g$-gon and rotate by $\frac{2\pi}{4g}$.

Reducible. A mapping class is said to be *reducible* if it fixes, set-wise up to homotopy, a multicurve, i.e. a disjoint union of simple closed curves. A reducible mapping class could be finite order. A Dehn twist is reducible and infinite order. In general, one can start with two surfaces, each with a single boundary component, and find a complicated mapping class on each that fixes the boundary (eg pseudo-Anosov). Then glue together ot get a reducible mapping class of a closed surface.

Pseudo-Anosov. A *pseudo-Anosov* homeomorphism of a surface is defined in exactly the same way as an Anosov homeomorphism of the torus, except that the foliation can have singularities (of the prong-type exhibited by quadratic differentials).

One construction of pseudo-Anosovs is via branched covers of tori with Anosov homeomorphisms. E.g. for a square-tiled surface, any hyperbolic matrix in $SL_2(\mathbb{Z})$ has a power that induces a pseudo-Anosov homeomorphism of the surface.

Theorem 5.3 (Nielsen-Thurston classification)**.** *Every mapping class $f \in Mod(S_g)$ is either periodic, reducible, or pseudo-Anosov.*

5.3 Tools for proof

Note that $\mathcal{M}_g$ is non-compact: we can construct a sequence of surfaces with injectivity radius going to 0 (and injectivity radius is a continuous, positive function on $\mathcal{M}_g$). The result below means that this is the only source of non-compactness.

Let $\ell(X)$ denote the length of the shortest closed geodesic on X.

Proposition 5.4 (Mumford compactness criterion)**.** *Let*

$$\mathcal{M}_g(\epsilon) := \{X \in \mathcal{M}_g : \ell(X) \geq \epsilon\}.$$

Then $\mathcal{M}_g(\epsilon)$ is compact.

Proof. Note first that $\mathcal{M}_g(\epsilon)$ is closed, since ℓ is a continuous function on $\mathcal{M}_g$.

It thus suffices to show that $\mathcal{M}_g(\epsilon)$ is contained in a compact set. We will show that this can be taken to be a finite union of compact sets which each correspond to a different topological type of Bers pants decomposition (see Proposition below). Note first that there are only finitely many different types of topological pants decomposition $\mathcal{P}$ on a fixed surface. For the first cuff curve the choices correspond to topological types of the complementary surface (g choices). Similarly, there are finitely many choices for the second curve, etc.

Now for each $\mathcal{P}$, we get a map $\mathbb{R}_{>0}^{3g-3} \times \mathbb{R}^{3g-3} \to \mathcal{M}_g$ from the Fenchel-Nielsen coordinates wrt $\mathcal{P}$ (the map is $\pi \circ FN^{-1}$, where $\pi : \mathcal{T}_g \to \mathcal{M}_g$ is the natural forgetful map, and FN is the Fenchel-Nielsen coordinate map we defined). Let $S_{\mathcal{P}} \subset \mathcal{M}_g$ be the image under this map of the set

$$\{(\ell_1, \dots, \ell_{3g-3}, \tau_1, \dots, \tau_{3g-3} : \epsilon \leq \ell_i \leq B_g, 0 \leq \tau_i \leq \ell_i \text{ for } i = 1, \dots, 3g-3\}.$$

By the Bers constant proposition, $\mathcal{M}_g(\epsilon)$ is contained in $\bigcup_{\mathcal{P}} S_{\mathcal{P}}$. Each $S_{\mathcal{P}}$ is compact, since it is the image of a compact set under a continuous map.

■

Proposition 5.5 (Bers constant)**.** *For each $g \geq 2$, there exists a constant B_g such that any $X \in \mathcal{M}_g$ has a pants decomposition with all cuff curves having length $\leq B_g$.*

Proof. Recall that by Gauss-Bonnet, the area of X is $\pi(4g-4)$. If we pick any point p on X, and grow a ball centered at p, for some radius R it must eventually be non-embedded, since otherwise the area of the ball would eventually exceed the area of X. This R can be made to depend only on g. From two radii of this ball, we can form a essential closed curve of length $\leq 2R$. We can take some piece of this and tighten to get a simple closed curve of length at most $2R$.

We then continue in a similar manner in the complement of the surface (this is actually a bit subtle, since the new surface we're working with has boundary). ■

Lemma 5.6 (Collar)**.** *Let γ be a simple closed curve on a closed hyperbolic surface X. The set*

$$C(\gamma) := \{x \in X : d(x, \gamma) \leq R\}$$

is an embedded annulus, where $R = \sinh^{-1}\left(\frac{1}{\sinh(\ell(\gamma)/2)}\right)$.

Note that R is roughly $\log(1/\ell)$.

Proof. We will prove the weaker/softer result that there is a collar of width going to infinity as $\ell(\gamma) \to 0$.

It suffices to show that for a right-angled hexagon with a short side a, the neighboring sides B and C are long. We can assume that the corner Ba is at the point i in upper half space, and side B points down from i. Since a is short, each of the geodesics containing B, C must have an endpoint near 0. Now side A must be ultraparallel to both B, C (since there exists an orthogeodesic to each). It follows that that the endpoints of geodesic A must be sandwiched between the two points near 0 described above. But then both B must be long, since otherwise c would not intersect A. By the same argument, C is also long. ■

Corollary 5.7. *There is a constant $\delta > 0$ such that any two simple closed geodesics of length less than δ on a hyperbolic surface are disjoint.*

Corollary 5.8. *There is a contant δ such that on any $X \in \mathcal{T}_g$, there are at most $3g-3$ simple closed geodesics of length less than δ.*

Lemma 5.9 (Discreteness of length spectrum)**.** *Let X be a closed hyperbolic surface. Given $L > 0$, there are finitely many closed geodesics on X of length less than L.*

Proof. Consider a (compact) fundamental domain D for the action of $\pi(X)$ on $\mathbb{H}$, and pick a point $p \in D$. By proper discontinuity of the action, only finitely many images of D under $\pi(X)$ intersect $B(p, L)$. These correspond to the only closed geodesics that can have length at most L. ∎

Proposition 5.10 (Proper discontinuity)**.** *The action of $Mod(S_g)$ on $\mathcal{T}_g$ is* properly discontinuous*: for each $K \subset \mathcal{T}_g$ compact, the set*

$$\{g \in Mod(S_g) : gK \cap K \neq \emptyset\}.$$

is finite.

It follows that $\mathcal{M}_g = \mathcal{T}_g/Mod(S_g)$ is an *orbifold*.

Proof. The first fact that we will use is: given a compact set $K \subset \mathcal{T}_g$ and constant L, there are only finitely many isotopy classes of curves that have length less than L on some $X \in K$. This follows from Discreteness of length spectrum, together with the fact that all lengths of curves change by a bounded amount in a Teichmüller ball (Wolpert's lemma).

The second fact is: we can find a collection of closed curves $\mathcal{S}$ on S_g, such that a mapping class f is determined, up to finite ambiguity, by the action of f on $\mathcal{S}$.

Now we put these two facts together. Suppose the proposition were false. Then we can find a compact K and distinct $g_1, g_2, \ldots$ with $g_i K \cap K \ni X_i$, for some X_i. By the second fact, there must be infinitely many sets among the $g_i(\mathcal{S})$. Now by continuity we can find a C such that any element of $\mathcal{S}$ has length at most C on all surfaces in K, and in particular on every X_i. Since the action of the mapping class group just remarks curves, all the curves in $g_i(\mathcal{S})$ must have length at most C on X_i. This contracts the first fact, since $\cup_i g_i(\mathcal{S})$ is infinite. ∎

Proposition 5.11 (Finiteness of isometry group)**.** *The orientation preserving isometry group $Isom^+(X)$ of any hyperbolic surface is finite.*

Once we know it's finite we get that the quotient $X/Isom(X)$ is compact hyperbolic orbifold. One can classify these, and use this to show the uniform bound $|Isom(X)| \leq 84(g-1)$.

Proof. First we show that $Isom^+(X)$ is a compact topological group. This follows from an Arzela-Ascoli argument (using only that we have a compact Riemannian manifold).

Next, we show that $Isom^+(X)$ is discrete. It suffices to show that any two isotopic isometries are equal; since we are working in a group, we can assume that one of the isometries is the identity. Lifting the other isometry to the universal cover $\mathbb{H}$, we see that it moves points a bounded amount. But then, by the classification of hyperbolic isometries, it must be the identity.

Since a compact, discrete space is finite, we are done. ■

5.4 Proof of Nielsen-Thurston classification

Proof of Theorem 5.3. We will break up into cases based on translation length τ.

Case 1: $\tau = 0$ and achieved.

If $X \in \mathcal{T}_g$ is fixed by f, then f must act by an isometry. By Finitness of isometry group, f must be periodic.

Case 2: τ is not achieved.

Let X_i be a sequence with $d(X_i, fX_i) \to \tau(f)$. We first claim that $\pi(X_i) \to \infty$, where $\pi : \mathcal{T}_g \to \mathcal{M}_g$ is the natural projection. Suppose the contrary. Then we can find h_i mapping classes, $X_i \in \mathcal{T}_g$, and $Y \in \mathcal{T}_g$ such that $h_i X_i \to Y$ (after passing to a subsequence). By invariance of the Teichmüller metric under mapping class group action, we have

$$d(h_i X_i, h_i f h_i^{-1}(h_i X_i)) = d(h_i X_i, h_i f X_i) = d(X_i, fX_i) \to \tau(f).$$

It then follows from the triangle inequality that

$$d(Y, h_i f h_i^{-1} Y) \leq d(Y, h_i X_i) + d(h_i X_i, h_i f h_i^{-1} h_i X_i) + d(h_i f h_i^{-1} h_i X_i, h_i f h_i^{-1} Y).$$

The first and last terms are equal, and go to zero, and the middle term's limiting behavior was found above, so we get $d(Y, h_i f h_i^{-1} Y) \to \tau(f)$.

Now by Proper discontinuity, $h_i f h_i$ must equal some fixed map g along a subsequence. But then $d(Y, gY) = \tau(f)$, contradiction.

It follows from the Mumford compactness criterion that for any $\epsilon > 0$, for large enough i, all X_i have an ϵ-short simple geodesic. By Wolpert's lemma, the length of such a short geodesic on fX_i is at most $\exp(2\tau f)\epsilon$. Now, for ϵ sufficiently small, there are most $3g-3$ simple closed geodesics of length at most ϵ (Corollary 5.8). Making ϵ very small, we get that f must act on a fixed finite set of disjoint simple closed curves, i.e. f is reducible.

Case 3: $\tau > 0$ and achieved.

Let $X \in \mathcal{T}_g$ achieve the translation distance, i.e. $d(X, fX) = \tau$. Let γ be the (unique) Teichmüller geodesic joining X, fX. We first show that $f(\gamma) = \gamma$. Suppose $Y \in \gamma$ is (strictly) between X and fX. We wish to show $fY \in \gamma$. By definition of translation length, $d(Y, fY) \geq \tau$. On the other hand, there is a path α from Y to fY by first going along γ from Y to fX, and then following the (unique) geodesic segment connecting fX, fY, which has length $d(fX, fY) = d(X, Y)$. Thus

$$|\alpha| = d(Y, fX) + d(fX, fY) = d(Y, fX) + d(X, Y) = d(X, fX) = \tau,$$

i.e. α is an optimal path from Y to fY. By Teichmüller's uniqueness theorem, α must be a Teichmüller geodesic. Since the first part of α is along γ, it follows that α itself must be a subsegment of γ. In particular, fY is on γ. By applying this argument repeatedly on longer and longer subsegments of γ, we get that $f(\gamma) = \gamma$.

We then get that the terminal quadratic differential associated to the Teichmüller map from X to fX along γ equals the initial differential for the Teichmüller map from fX to f^2X along γ. It follows that f is pseudo-Anosov. ∎

6 Deligne-Mumford compactification

Goal: compactify $\mathcal{M}_g$ in a "natural" way.

Augmented Teichmüller space. A sequence of Riemann surfaces X_ϵ with systoles of length ϵ going to 0 (e.g. by fixing a pants decomposition, and sending one Fenchel-Nielsen length coordinate ℓ to zero, while keeping all other parameters fixed) escapes every compact set of $\mathcal{M}_g$. We can think of the limiting object as a Riemann surface where $\ell = 0$; there is no longer a twist parameter.

Attaching all such limits, gives the augmented Teichmüller space $\overline{\mathcal{T}_g}$. This is a somewhat nasty space: it is not even locally compact, since neighborhoods of boundary points contain surfaces with arbitrarily large twist parameters. However, the mapping class group acts on $\overline{\mathcal{T}_g}$, and the quotient $\overline{\mathcal{M}_g}$, the Deligne-Mumford compactification, is a compact, smooth orbifold. (In fact, it also has a nice algebraic/analytic structure; it is a projective variety).

Model for the singularity. The objects we are adding are nodal algebraic curves. Near the node, the curves look like the locus $wz = \epsilon$ in $\mathbb{C}^2$.

Ends of moduli space. A topological space Y is said to have *one end* if for every compact set $K \subset Y$, we have that $Y - K$ has one component whose closure is non-compact.

Theorem 6.1. *Moduli space $\mathcal{M}_g$ has one end.*

Proof. Suppose the contrary for some K. The complement certainly has at least one component whose closure is non-compact. Using the below proposition we can find an open set $U \subset \overline{\mathcal{M}_g}$ such that

- U is disjoint from K,
- $U \supset \partial\mathcal{M}_g$,

- $U - \overline{\mathcal{M}_g}$ is connected (using that $\overline{\mathcal{M}_g}$ is a smooth complex orbifold).

By the compactness of $\overline{\mathcal{M}_g}$, we have that $\overline{\mathcal{M}_g} - U = \mathcal{M}_g - U$ is compact. Hence any components of $\mathcal{M}_g - K$ with non-compact closure intersect U. But since U is connected and disjoint from K, there can only be one such component.

■

Proposition 6.2. *The boundary $\partial\mathcal{M}_g := \overline{\mathcal{M}_g} - \mathcal{M}_g$ is connected.*

Proof. Let γ be a simple closed curve on S_g; there are only finitely many choices up to action of the mapping class group. Let V_γ be the locus where γ has been pinched, and no other curves.

First note that $\partial\mathcal{M}_g$ is the union of $\overline{V_\gamma}$ over the finitely many choices of γ.

Secondly, V_γ is itself a product of moduli spaces of Riemann surfaces with marked points, and is hence connected; thus so is $\overline{V_\gamma}$.

Thirdly, $\overline{V_\gamma}$ and $\overline{V_{\gamma'}}$ intersect for any γ, γ': this is because we can find representatives of the mapping class group orbits of the curves that are disjoint from one another, which means that V_γ and $V_{\gamma'}$ admit a common degeneration.

These three facts together imply $\partial\mathcal{M}_g$ is connected. ■

7 Weil-Petersson metric

There is a Hermitian form on the tangent bundle $T^*\mathcal{T}_g \cong QD(X)$, given by

$$\langle q_1, q_2 \rangle := \int_X q_1 \overline{q_2} ds^{-2},$$

where ds is the hyperbolic metric on X. This induces a dual metric on the tangent space, called the *Weil-Petersson form* on Teichmüller space. It descends to a form on moduli space. It is a Kähler form, which means that it induces both a Riemannian metric and a symplectic form.

Theorem 7.1 (Wolpert). *The Weil-Petersson symplectic form ω_{WP} is equal to*

$$\sum_i d\ell_i \wedge d\tau_i,$$

where the sum ranges over the Fenchel-Nielsen coordinates with respect to any pants decomposition.

One interesting consequence of the formula is that the expression above is independent of the choice of pants decomposition.

Theorem 7.2. *$\mathcal{M}_g$ is an incomplete metric space with respect to the Weil-Petersson Riemannian metric. Its metric completion is homeomorphic to the Deligne-Mumford compactification $\overline{\mathcal{M}_g}$.*