Math 851 Notes
Teichmüller Theory

Lectures by Autumn Kent
Notes by Daniel Hast

Fall 2014

Contents

1 2014-09-03: Moduli problems

Given a geometric problem, it's often better or easier to consider a *family* of objects instead of a single object.

You've seen "moduli spaces" before: e.g., $\mathbb{P}^n$ is the "moduli space of lines through the origin" in $\mathbb{A}^{n+1}$. Likewise, Grassmannian varieties are moduli spaces of linear subspaces.

1.1 Example: Polygonal billiards

Say that we're interested in billiards on a polygon. Given a billiard table, is there a periodic orbit? For rectangles, there is. The question is more interesting for nonstandard tables.

Question ($\tilde{2}00$ years old): Does every triangle has a periodic orbit? If the triangle is acute, then there is (Fagnano, 1775). For right triangles, there's a periodic orbit along the same lines as for rectangles.

Theorem 1.1 (Schwartz, 2008)**.** *There's a periodic billiards orbit on any triangle with largest angle* $\leq 100°$*.*

Idea: Start with a particular triangle you understand that has a periodic orbit. Perturb and hope for stability. More precisely, look at some moduli space of all triangles, and try to cover it with sets of triangles with periodic orbits.

However:

Theorem 1.2 (Hooper)**.** *This doesn't quite work, because every periodic orbit in a right triangle is unstable.*

Theorem 1.3 (Masur, 1986)**.** *Any rational polygon has a periodic orbit.*

The proof of Masur's theorem uses the full force of Teichmüller theory (moduli of Riemann surfaces of high genus).

1.2 Moduli of Riemann surfaces

We're interested in moduli of Riemann surfaces.

Remark 1.4. One can view geodesic closed loops on Riemann surfaces as arising from periodic billiard orbits.

Let's warm up to studying moduli (or "deformations") of surfaces.

Example 1.5. Consider the "flat" torus: locally, it looks like $\mathbb{R}^2$, both topologically and geometrically. (It has a flat Riemannian metric.) One can construct flat tori by identifying opposite sides of a parallelogram.

Let's study the "space" of flat tori, up to isometry. We could always scale the plane, which is boring; to get rid of this problem, we'll assume our tori all have area 1.

Theorem 1.6. *Every flat torus is obtained from a parallelogram by identifying opposite sides.*

Proof. Take two curves along the torus that intersect exactly once. By a compactness argument, we can shorten the curves to be locally straight.

Lemma. *The straightened loops still intersect exactly once.*

Then cut along those curves to obtain a parallelogram. □

Now we want to figure out the moduli space of parallelograms. Next time, we'll go back and more carefully parametrize parallelograms.

2 2014-09-05: Deformations of tori

Consider a flat torus T (of area 1) given by a parallelogram with opposite sides identified. Abstractly, this is a topological torus with a collection of charts $T \supset U \xrightarrow{\simeq} V \subset \mathbb{R}^2$ where the overlaps are isometries.

What's the space of all such flat tori? As a set, we're interested in flat tori up to isometry. How do we parametrize this space? If we just start listing tori, then there's redundancy.

Let's start again and try to parametrize flat tori in a less naive way:

- Instead of normalizing to area 1, we'll assume the base lies horizontally and has length 1. The upper-right vertex $x + iy \in \mathbb{C}$ in the upper-half plane now determines the parallelogram.
- By slicing parallelograms diagonally, we see another source of redundancy. Thus, we can restrict to any vertical strip of width 1, say $\left[\frac{1}{2}, \frac{3}{2}\right]$.
- Inversion in the circle of radius 1 centered at $1 \in \mathbb{C}$ preserves equivalence classes.
- Reflection across the vertical line $y = 1$ also preserves equivalence classes.

So, now we parametrize flat tori by points in the region

$$\left\{ z \in \mathbb{C} \;\middle|\; \tfrac{1}{2} \leq \mathrm{Re}(z) \leq 1,\ |z - 1| \geq 1 \right\}.$$

We claim that every flat torus corresponds to a *unique* point in the above region.

3 2014-09-08: Points at infinity

Recall the fundamental domain corresponding to the moduli space of flat tori (up to isomorphism and scaling).

If we just start going up in the fundamental domain (with flat tori normalized to area 1), there's a preferred loop that starts to get short. So, at ∞, this loop has length zero. So, ∞ "is" that loop! (Likewise, as we go down, we see a different curve get short.)

Moving around, we see certain points "at infinity" correspond to homotopy classes of loops that are getting short.

When we were finding our fundamental domain, at each stage, we simplified the region in some nice way, and each one corresponded to a homeomorphism of the region. For example, there's inversion and reflection. The homeomorphism corresponding to translation is a "Dehn twist": cut open along the "base" and twist by 2π.

In other words, the upper-half plane $\mathbb{H}$ is a space of parallelograms, and we have a group G of (homotopy classes of) homeomorphisms acting on $\mathbb{H}$, and our moduli space is $\mathbb{H}/G$. We also have points at ∞ corresponding to loops getting short.

4 2014-09-10: The upper-half plane as a space of metrics

Given a parallelogram P determined by a point $\omega \in \mathbb{H}$, there is a unique affine map from the unit square to P that preserves edge labels and orientations. Now, pullback the Euclidean metric on $\mathbb{C}$ along this map and scale to give the unit square area 1. Call that metric X_ω.

For any (p, q)-curve γ_{pq} (with $\gcd(p, q) = 1$), we can make γ_{pq} short by shearing until γ_{pq} is vertical, then take $\operatorname{Im}(\omega) \to 0$ along a vertical line, and the length $\ell_{\omega_s}(\gamma_{pq}) \to 0$.

Now, for each (p, q) with $\gcd(p, q) = 1$, we have a point in $\mathbb{R} \cup \{\infty\} = \partial\mathbb{H}$.

Nathan's question: If we head to $\mathbb{R}$, does some curve necessarily get short? The answer: No! Let γ be a "foliation" of the flat square T by lines of an irrational slope.[1] We can't say that γ gets short when $\operatorname{Im}(\omega) \to 0$.

Given a closed curve δ, there is a geodesic representative δ_* with respect to X_ω. So δ_* and γ meet at some angle in X_ω.

5 2014-09-12: Geodesics in the upper-half plane

We saw we can shrink any curve, and each curve gave us a point in $\mathbb{R} \cup \{\infty\} = \partial\mathbb{H}$. Also, we found points in $\partial\mathbb{H}$ where no curve is getting short (in fact, every curve gets long). Think of $\partial\mathbb{H}$ as a circle's worth of foliations.

Starting from a square torus, as we headed toward some $F \in \partial\mathbb{H}$, the metric on the torus contracts along the corresponding foliation $\mathcal{F}$ and expands in the perpendicular direction. (At least, this is happening at the *end* of our path.)

Is this the most "efficient" way to shrink $\mathcal{F}$? No — we can "go up" first! If we let ω_n be the metric on the rectangular torus with height $\sqrt{n}$ and width $\frac{1}{\sqrt{n}}$, then

$$\ell_{X_{\omega_n}}(\gamma_{-n,1}) = \sqrt{2n} \ll n.$$

Going up made the length drop. However, eventually this effect decays.[2] Moving up and to the right and then down again seems like a good way to shrink $\gamma_{-n,1}$.

It turns out that the most efficient path is to follow a half-circle centered on the real line, i.e., moving along geodesics in the hyperbolic plane, $\mathbb{H}$ with the metric

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

The isometries $\operatorname{Isom}(\mathbb{H})$ of this space are Möbius transformations, and

$$\operatorname{Isom}^+(\mathbb{H}) = \operatorname{PSL}_2(\mathbb{R}).$$

The distance in $\mathbb{H}$ between $ai + c$ and $bi + c$ (with $b > a$) is $\log \frac{b}{a}$.

Observation: $\log b$ is the "eccentricity" of the torus X_{bi+1} as compared to X_{i+1}.

[1]"It's called a foliation because it's like a book, but you can never open it."

[2]Like Icarus' wings as he approached the Sun.

6 2014-09-15

Think of the upper-half plane as a space of flat tori of area 1. There is a group Γ acting on $\mathbb{H}$ that gives us area-1 flat tori *up to isometry*. Elements of the group correspond to homeomorphisms.

We can give $\mathbb{H}$ the topology from $\mathbb{C}$, and even better, the hyperbolic metric $ds^2 = \frac{dx^2+dy^2}{y^2}$. The group $\Gamma^+ = \mathrm{PSL}_2(\mathbb{Z})$ acts by fractional linear transformations on $\mathbb{H}$ by isometries, and we can consider the quotient spaces $\mathbb{H}/\Gamma$ or $\mathbb{H}/\Gamma^+$. This gives a moduli space $\mathcal{M}(T)$, which is an orbifold.

Exercise 6.1. Think about $\mathrm{Stab}_{\Gamma^+}(X_{e^{\pi i/3}})$. Hint: $T^2 = \text{hexagon}/(\text{side pairings})$. Think about self-isometries.

Geodesics in $\mathbb{H}$ correspond to efficiently "shrinking" the metric, i.e., locally stretching the flat metric. We can make this local stretching into a metric on $\mathbb{H}$ by saying the distance is log(amount of stretching). For example, local stretching of the form

$$\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$$

has distance t.

6.1 Higher genus

We want to do all this in higher genus; however, there are no flat metrics on a surface of genus $g \geq 2$. But what we really want to classify are Riemann surfaces with complex structure. (In the genus 1 case, uniformization shows that area-1 flat tori are equivalent to Riemann tori, i.e., tori with complex structure and holomorphic transition maps.)

7 2014-09-17: Higher genus surfaces

We want to study Riemann surfaces of arbitrary genus.

A *Riemann surface* is:

(1) a topological surface S of finite type, i.e., a compact surface with finitely many points removed[3] ("punctures");

(2) a collection of charts on S with holomorphic transition maps.

Morphisms of Riemann surfaces are locally holomorphic smooth maps.

Moduli problem: Given S, what are all the Riemann surfaces homeomorphic to S?

Example 7.1. The Riemann sphere $\hat{\mathbb{C}}$ is the only Riemann surface homeomorphic to S^2.

Example 7.2. The complex plane $\mathbb{C}$ is *almost* the only Riemann surface homeomorphic to $\mathbb{R}^2$; the only other one is the unit disk Δ.

[3] Caveat: For us, we usually require that a "puncture" is $\Delta - \{0\}$ (where Δ is the open unit disk) and not $\Delta - \overline{\frac{1}{2}\Delta}$.

The Riemann mapping theorem says that $\hat{\mathbb{C}}$, Δ, and $\mathbb{C}$ are the only simply-connected Riemann surfaces.

If we puncture $\mathbb{C}$, we get $\mathbb{C} - \{w\}$. They are all equivalent: translations are holomorphic and act transitively on $\mathbb{C}$. It has no moduli (when considered as a "punctured" surface).

If we puncture again, we get $\mathbb{C} - \{v, w\} \cong \mathbb{C} - \{0, 1\}$, a thrice-punctured sphere. Note that $\mathrm{PSL}_2(\mathbb{C})$ acts transitively on triples of distinct points in $\hat{\mathbb{C}}$.

However, when we puncture four times, we get $\mathbb{C} - \{0, 1, w\}$, and these are different for different w. This gives us a moduli space of four-times-punctured spheres, corresponding to points of $\mathbb{C} - \{0, 1\}$.

Puncturing $\hat{\mathbb{C}}$ five times, we get $\mathbb{C} - \{0, 1, w, w'\}$ with $w \neq w'$; this gives us a 2-dimensional moduli space.

8 2014-09-19: More about Riemann surfaces

Continuing from last time, let $S_{0,n}$ be the 2-sphere with n punctures. Let $\mathcal{M}(S_{0,n})$ be the moduli space of (finite-type) Riemann surfaces homeomorphic to $S_{0,n}$. This is a smooth complex manifold. Think of it as a "configuration space" of points in $\hat{\mathbb{C}}$.

9 2014-09-22: Higher genus, continued

How do we find families of Riemann surfaces homeomorphic to S, i.e., how to find families of complex structures on S?

We looked at $S_{0,n}$ earlier: $S_{0,n}$ seems to have a space of $\mathbb{C}$-structures of complex dimension $n - 3$. (This was known to Riemann in a vague sense.)

For higher genus, let $S_{g,n}$ be the surface of genus g with n punctures. Given a Riemann surface X, we can start deleting points, and you will (eventually) get "new" deformations (aside from some "accidents", e.g., $\mathcal{M}(S_{1,0})$ and $\mathcal{M}(S_{1,1})$ both have $\mathbb{C}$-dimension 1).

What about $S_{g,0}$ for $g > 1$? Let's focus on $S = S_{2,0}$. Think about how we might find $\mathbb{C}$-structures.

Problem: S admits no flat metric. One way to see this is to look at the curvature. Assume we know that Riemannian manifolds live in some $\mathbb{R}^n$. We then think of curvature as the determinant of the Gauss map. By Gauss–Bonnet,

$$\int_S \kappa(x) = 2\pi\chi(S),$$

so $\chi(S_{g,0}) = 2 - 2g \neq 0$. Hence, there are no flat metrics on $S_{g,0}$ for $g > 1$.

Instead, we use *singular* flat metrics. For example, for $g = 2$, take a regular octagon and identify opposite sides via *translations*. We get a genus 2 surface X that looks like there's a flat metric (inherited from $\mathbb{R}^2$), but "too much angle" at a point.

Definition 9.1. A *singular flat metric* is a metric that's locally isometric to $\mathbb{R}^2$ except at a finite set, where the angle is an integer multiple of π.

Affine transformations give us new *octagons*. We get a family of singular flat metrics, and hence a family of complex structures.

10 2014-09-24

10.1 Comparison of geometries

Model geometries: $\mathbb{H}_k$ ($k < 0$), $\mathbb{S}_k$ ($k > 0$), and $\mathbb{E}$ ($k = 0$). There are CAT(k) spaces.

10.2 Singular flat metrics, continued

Singular flat metrics are flat metrics at all but a finite set of points, where we have a "cone point" with some angle $\alpha\pi$.

Consider again the octagon with opposite sides identified by translations, giving a surface S of genus 2. We can obtain new $\mathbb{C}$-structures on S by deforming this octagon, e.g.: Stretch this octagon with some affine map, get some new flat metric, and hence some new $\mathbb{C}$-structure.

It would be naive to expect this to give we all $\mathbb{C}$-structures. For example, what about non-convex octagons?

Theorem 10.1. *Look at this family where we start with the regular octagon and do the simplest thing with affine stretches. Then this family is naturally the hyperbolic plane* $\mathbb{H}$.

Moreover (Veech), if we take $\mathbb{H}/(\textit{isomorphism of } \mathbb{C}\textit{-structure})$, *we get a finite-volume hyperbolic orbifold.*

z is a nice coordinate patch, but it isn't really a well-defined thing on S. On the other hand, dz is a nice 1-form on S. There is a 1-form $\varphi(z)\,dz$ on S that looks like dz everywhere except at p, where it looks like $z^2\,dz$.

We can now recover the flat metric: integrating $\varphi(z)\,dz$ gives a coordinate

$$\zeta(w) = \int_0^w \varphi(z)\,dz,$$

called the natural parameter. There are well-defined horizontal and vertical directions, so we pull back the flat metric from $\mathbb{C}$.

11 2014-09-26

We found a 1-form φ that looks like dz away from p, and at p, it looks like $z^2\,dz$. Hence, a flat metric induces a $\mathbb{C}$-structure and a holomorphic 1-form φ.

Now, given a $\mathbb{C}$-structure, i.e., a Riemann surface $X \cong S$ and a holomorphic 1-form φ, we obtain a flat metric by integration:

$$\zeta(w) = \int_0^w \varphi(z)\,dz$$

is a "coordinate", called the *natural parameter* (actually a coordinate away from zeros of φ). At a zero, it's a "branched" coordinate.

We can use ζ to pull back the flat metric on $\mathbb{C}$ to get a singular flat metric on X. Not only that, but we can pull back the "horizontal" and "vertical" dilations. And, moreover, these horizontal and vertical foliations are "orientable".

12 2014-09-29

Given a Riemann surface X, a holomorphic quadratic differential q on X is a holomorphic section of the (symmetric) square of the holomorphic cotangent bundle. (Given X, there are a lot of these. They form a vector space of $\mathbb{C}$-dimension $3g-3$ when $X \cong S_{g,0}$.)

What's the Teichmüller space of S? It's supposed to capture moduli of S, but intelligently parametrized.

As a set, what is it? Fix a topological oriented surface S. A *marked* Riemann surface is a pair (X, f) where X is a Riemann surface and $f : X \to S$ is a homeomorphism. Two marked Riemann surfaces (X, f) and (Y, g) are *Teichmüller-equivalent* if there is an isomorphism of Riemann surfaces (i.e., a biholomorphic homeomorphism) $\eta : X \to Y$ such that η is isotopic to $g \circ f^{-1}$.

Proposition 12.1 (Alexander's trick)**.** *Let D^2 be a closed disk, and let $f : D^2 \to D^2$ be a continuous map such that $f|_{\partial D^2} = \mathrm{id}$. Then f is isotopic to the identity map.*

Shrink all of the "bad" part of the map to a point. This is related to the "bachelor's unknotting": maps $S^1 \to S^3$ up to isotopy in the topological category are trivial.

Theorem 12.2 (Baer–Epstein)**.** *Given closed surfaces X, Y, two homeomorphisms $f, g : X \to Y$ are homotopic if and only if they're isotopic.*

The Teichmüller space for fixed S is $\{(X, f)\}$ /(Teichmüller equivalence). To get a topology, look at the map $g \circ f^{-1}$ and measure the deviation from being holomorphic.

Biholomorphic means *conformal* ("angle-preserving"); infinitesimally, it means the derivative preserves circles. Given a diffeomorphism, infinitesimal S^1's go to ellipses, and the "eccentricity" measures deviation from being conformal.

13 2014-10-01: Teichmüller space

Let S be a closed oriented surface. One way to move around in Teichmüller space $\mathcal{T}(S)$ would be this: Given $f : S \to X$, we could simply change f to a different homeomorphism. This gives some potentially large set of points in $\mathcal{T}(S)$.

For example, consider X with a flat metric, cut along a geodesic loop γ, twist $360°$, and glue back together.

This usually won't usually give us the same point, but it could. If the homeomorphism induces an isomorphism of X, then you'll get nowhere in $\mathcal{T}(S)$.

Aside 13.1. Alternative to marking is to think of $\mathcal{T}(S)$ as being a set of atlases on S up to isomorphisms isotopic to id_S.

This is a space?

Given a quadratic differential q on X, we can deform the $\mathbb{C}$-structure by using a natural parameter ζ and applying a Teichmüller deformation $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$.

If we fix a marking, this changes the Teichmüller class. The distance is supposed to be t.

14 2014-10-03

Understanding the moduli space $\mathcal{M}(S) = \{\text{Riemann surfaces homeomorphic to } S\}/(\text{isomorphism})$ is hard, since the parametrization is often redundant. So, instead, we study Teichmüller space $\mathcal{T}(S)$, the moduli space of *marked* Riemann surfaces homeomorphic to S, which maps to $\mathcal{M}(S)$ by forgetting markings.

The space $\operatorname{Diff}_+(S)$ of orientation-preserving diffeomorphisms up to isotopy "acts" on $\mathcal{T}(S)$, and this action factors through

$$\operatorname{Diff}_+(S)/\operatorname{Diff}_0(S) \cong \operatorname{Mod}(S),$$

where $\operatorname{Diff}_0(S)$ is the space of diffeomorphisms isotopic to id_S and $\operatorname{Mod}(S)$ is the mapping class group, defined by

$$\operatorname{Mod}(S) = \pi_0 \operatorname{Diff}_+(S) = \pi_0 \operatorname{Homeo}_+(S),$$

where the latter equality is a theorem. We have

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

Example 14.1. $\mathcal{T}(T^2) = \mathbb{H}$ is the space of marked flat metrics on $T^2 = S^1 \times S^1$. This has an action of the modular group $\mathrm{PSL}_2(\mathbb{Z})$, and the quotient is $\mathcal{M}(S) = \mathrm{PSL}_2(\mathbb{Z})/\mathbb{H}$.

Theorem 14.2 (Fricke, Klein)**.** $\operatorname{Mod}(S)$ *acts properly discontinuously on* $\mathcal{T}(S)$.

Theorem 14.3. *For $g \geq 2$, $\mathcal{T}(S)$ is a $\mathbb{C}$-domain of $\mathbb{C}$-dimension $3g-3$, i.e., is homeomorphic to $\mathbb{R}^{6g-6}$.*

Theorem 14.4 (Royden)**.** *For $g \geq 2$,* $\operatorname{Mod}(S)$ *is the entire group of biholomorphic automorphism of $\mathcal{T}(S)$.*

What is the topology on $\mathcal{T}(S)$? We'd like to measure the failure of conformality of a homeomorphism $f : X \to Y$. We could locally measure the eccentricity of ellipses, and take the distance to be the infimum over homotopy classes of f of $\frac{1}{2}\log(\text{maximum eccentricity of } f)$.

Let $f : \Omega \to \mathbb{C}$ be a C^1 embedding from a domain $\Omega \subseteq \mathbb{C}$. From the perspective of complex analysis, f being conformal is the same as $f_{\overline{z}} = 0$, where

$$f_z = \frac{1}{2}\left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right) f,$$
$$f_{\overline{z}} = \frac{1}{2}\left(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\right) f.$$

One can check that, for $z \in \Omega$, the eccentricity $K_f(z)$ is given by

$$K_f(z) = \frac{|f_z| + |f_{\overline{z}}|}{|f_z| - |f_{\overline{z}}|}.$$

In general, we don't have $f_{\overline{z}} = 0$, but $f_{\overline{z}} = \mu(z) f_z$ for some μ with $\|\mu\|_\infty < 1$. Define

$$K(f) := \|K_f(z)\|_\infty,$$

the *dilatation* of f.

Given a diagram

$$\begin{tikzcd} S \arrow[r, "m_1"] \arrow[dr, "m_2"'] & X \arrow[d, "f"] \\ & Y \end{tikzcd}$$

with f a C^1-diffeomorphism, we define the *Teichmüller metric* by

$$d_{\mathcal{T}(S)}\big((X, m_1), (Y, m_2)\big) = \inf_{\substack{\text{homotopy} \\ \text{classes of } f}} \tfrac{1}{2}\log(K(f)).$$

Definition 14.5 (Regularity)**.** A homeomorphism $f : X \to Y$ is K-*quasiconformal* if:

(1) f has locally integrable distributional derivatives $f_z, f_{\overline{z}}$.

(2) $|f_{\overline{z}}| \leq k\,|f_z|$, where $k = \frac{K-1}{K+1}$.

To explain the idea of distributional derivatives, recall integration by parts (the product rule):

$$uf = \int uf' + \int u'f.$$

If there's an $f' \in C^1$ satisfying this for all smooth test functions u, then f' is called the *distributional derivative* of f.

Theorem 14.6 (Teichmüller)**.** *The homotopy class of f contains a unique "extremal" quasiconformal map (realizing the distance), and this extremal map is a "Teichmüller map" (i.e., a Teichmüller deformation for some quadratic differential q on X).*

15 2014-10-06: Quasiconformal maps

Definition 15.1 (quasiconformal)**.** Let $K \geq 1$. A homeomorphism $f : X \to Y$ between Riemann surfaces is K-*quasiconformal* if f has locally integrable distributional derivatives f_z and $f_{\overline{z}}$, and $|f_{\overline{z}}| \leq k\,|f_z|$, where $k = \frac{K-1}{K+1}$.

Note that conformal is the same as 1-quasiconformal, and if $K' \geq K$, then all K-quasiconformal maps are K'-quasiconformal.

There's an equation $f_{\overline{z}} = \mu(z) f_z$ for some $\mu(z) \in L^\infty(X)$ with $\|\mu\|_\infty < 1$; μ is called the "*complex dilatation*".

The *dilatation* of f is the smallest K for which f is K-quasiconformal.

Given Riemann surfaces X and Y and a quasiconformal map $f : X \to Y$, we say that f is *extremal* if its dilatation is minimal in its homotopy class.

Theorem 15.2 (Teichmüller)**.** *Given marked Riemann surfaces $\alpha : S \to X$ and $\beta : S \to Y$, there is a unique extremal map $f : X \to Y$ homotopic to $\beta\alpha^{-1}$. Moreover, f comes from a Teichmüller deformation for some holomorphic quadratic differential q on X.*

The space of all holomorphic quadratic differentials on X is a vector space $Q(X)$. By the Riemann–Roch theorem, $Q(X)$ is finite-dimensional of complex dimension $3g-3$.

Let $Q^1(X)$ be the unit ball in $Q(X)$. Then

$$\mathcal{T}(S) \cong \text{cone on } Q^1(X) \cong \mathbb{R}^{6g-6}.$$

Not only is $\mathcal{T}(S) \cong \mathbb{R}^{6g-6}$, but $\mathcal{T}(S)$ is a smooth manifold whose cotangent space is $Q(X)$. The complex dilatation more properly measures the deformations of a $\mathbb{C}$-structure. The tangent space is

$$T_X \mathcal{T}(S) = \{\text{differentials of type } (-1,1)\},$$

called the space of *Beltrami differentials*. (Example: $\frac{d\bar{z}}{dz}$.)

16 2014-10-08: Teichmüller's theorem

Why would we hope for Teichmüller's theorem? Consider Grötzsch's problem for a rectangle: Let f be a C^1 homeomorphism between rectangles $(0, a, ib, a+ib)$ and $(0, a', ib', a'+ib')$ in the complex plane; assume f preserves vertices and their order.

There's an affine map $f_0 = \frac{a'}{a}x + i\frac{b'}{b}y$ between the same rectangles. Does f_0 minimize the dilatation, and is it the unique extremal f?

Theorem 16.1 (Grötzsch). *$K(f) \geq K(f_0)$, with equality iff $f = f_0$.*

Proof. First, $a' \leq \int_\alpha |f_x|\, dx$. Recall: $K_f(z) = \frac{|f_z|+|f_{\bar{z}}|}{|f_z|-|f_{\bar{z}}|}$, $|\langle x, y\rangle| \leq \|x\|\,\|y\|$ (with equality iff $x = cy$), and $J_f(z) = |f_z|^2 - |f_{\bar{z}}|^2$. Integrate

$$\begin{aligned} a'b &\leq \iint_R |f_x|\, dx\, dy = \iint_R |f_z + f_{\bar{z}}|\, dx\, dy \\ &\leq \iint_R |f_z| + |f_{\bar{z}}|\, dx\, dy = \iint_R K_f(z)^{1/2} \cdot J_f(z)^{1/2}\, dx\, dy \\ &\leq \sqrt{\iint_R K_f(z)\, dx\, dy \iint_R J_f(z)\, dx\, dy}. \end{aligned}$$

Hence,

$$(a'b)^2 \leq \iint_R K_f(z)\, dx\, dy \iint_R J_f(z)\, dx\, dy \leq abK(f)a'b'.$$

So $K(f) \geq \frac{a'b}{ab'} = \frac{a'/a}{b'/b} = K(f_0)$.

Now suppose we have equality. The first inequality being an equality implies that $f = u(x) + iv(y)$, so $J_f(z) = u_x v_y$. Also, $K_f(z) = \frac{u_x}{v_y}$. Equality in Cauchy–Schwarz implies $K_f(z) = cJ_f(z)$, so $K_f(z), J_f(z)$ are constant. Hence u_x, v_y are constant, so f is affine. $\square$

17 2014-10-10: Teichmüller's uniqueness theorem

Theorem 17.1 (Uniqueness)**.** *If $f_0 : X \to Y$ is a Teichmüller map (i.e., a quasiconformal map coming from a Teichmüller deformation corresponding to some quadratic differential q_X and some $t \in \mathbb{R}$) and $f : X \to Y$ is a quasiconformal map homotopic to f_0, then $K(f) \geq K(f_0)$, with equality if and only if $f = f_0$.*

Remark 17.2. Recall the inequality $e^t \operatorname{Area}(R) = e^t ab \leq \iint_R |f_X| \, dA$ from the proof of Grötzch's theorem, discussed last time. An analogue of this is key to the Teichmüller uniqueness theorem:

$$\sqrt{K_0} \operatorname{Area}_{q_X}(X) \leq \int_X |f_x| \, dA.$$

Once we have this, the argument is essentially the same, integrating over X instead of R and using Cauchy–Schwarz.

Setup: Let X, q_X, t be as above, and let q_Y be the quadratic differential giving the flat metric on Y.

For $p \in X$, let $\alpha_{p,L}$ be a horizontal arc in X through p, with length L on both sides of p. Consider

$$\ell_{q_Y}(f(\alpha_{p,L})) = \int_{-L}^{L} |f_x| \, dx \geq 2L\sqrt{K(f_0)} - M.$$

Hence,

$$(2L\sqrt{K(f_0)} - M) \cdot \operatorname{Area}(X) \leq \int_X \int_{-L}^{L} |f_x| \, dx \, dA = \int_{-L}^{L} \int_X |f_x| \, dA \, da = 2L \int_X |f_x| \, dA.$$

Dividing by $2L$ and taking $L \to \infty$ yields the desired inequality.

18 2014-10-13

Teichmüller's theorem is often stated as two theorems: the existence theorem that there's an extremal Teichmüller map, and the fact that it's unique. The proof of uniqueness is morally the same as the proof of Grötszch's theorem for rectangles.

There will be some quasiconformal map homotopic to f, namely, any smooth map in the homotopy class. (By compactness, smooth homeomorphisms on compact surfaces are quasiconformal.) We could try taking a sequence $f_n \simeq f$ with dilatations tending to $\inf_{g \simeq f} K(g)$, and try to find a limiting map.

What do we do instead? We'll think of $\mathcal{T}(S)$ from a different perspective: By uniformization, we can think of $\mathcal{T}(S)$ as a space of hyperbolic metrics on S.

The uniformization theorem says that, for $g \geq 2$, given a Riemann surface $X \cong S_{g,0}$, there is a discrete group of isometries Γ acting freely on $\mathbb{H}$, such that $X \cong \mathbb{H}/\Gamma$.

What are the isometries of $\mathbb{H}$? We'll seen all of the isometries: $\operatorname{Isom}(\mathbb{H})$ is generated by inversions in the circles perpendicular to $\mathbb{R}$, and $\operatorname{Isom}_+(\mathbb{H}) = \operatorname{PSL}_2(\mathbb{R})$.

19 2014-10-15: Hyperbolic surfaces

Previously, we talked about Teichmüller's uniqueness theorem. Existence is tougher.

We'd like the space of all Teichmüller deformations of X to be all of $\mathcal{T}(S)$, i.e., we want to hit all (marked) Riemann surfaces. To see this, we introduce hyperbolic geometry!

Theorem 19.1 (Uniformization). *Every Riemann surface Y of genus $g \geq 2$ is isomorphic (as a Riemann surface) to a hyperbolic surface.*

What's a hyperbolic surface? Let $S = S_{g,0}$ be a closed surface of genus $g \geq 2$. A *hyperbolic structure* on S is a homeomorphism $S \cong \mathbb{H}/\Gamma$, where Γ is a discrete torsion-free subgroup of $\mathrm{Isom}_+(\mathbb{H}) \cong \mathrm{PSL}_2(\mathbb{R})$. In other words, a hyperbolic surface is a Riemann surface (S, g) such that g has constant curvature -1.

Hence, the universal cover $(\tilde{S}, \tilde{g})$ of (S, g) with the pullback metric is isometric to $\mathbb{H}$, and the deck group is $\Gamma \cong \pi_1(S)$.

Lemma 19.2. $\mathbb{H}/\Gamma \cong S$.

Think of hyperbolic metrics on S as *faithful* representations

$$\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R})$$

with discrete image, called the *holonomy representation* of the hyperbolic structure.

Exercise 19.3. Show that T^2 doesn't admit a hyperbolic structure.

Given a holonomy representation $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R})$, how do we deform it? Try to perturb the generators.

Theorem 19.4 (Poincaré). *A "small" deformation of a holonomy representation $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R})$ produces a holonomy representation $\rho' : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R})$.*

The set of holonomy representations $\mathcal{H}(S)$ is naturally a topological space, with topology induced by the inclusion

$$\mathcal{H}(S) \subset \mathcal{R}(S) := \mathrm{Hom}(\pi_1(S), \mathrm{PSL}_2(\mathbb{R})) \subset \mathrm{PSL}_2(\mathbb{R})^{2g} = \mathrm{Hom}(F_g, \mathrm{PSL}_2(\mathbb{R})).$$

Theorem 19.5. *$\mathcal{H}(S)$ is a connected component of $\mathcal{R}(S)$.*

What does $\mathcal{R}(S)$ look like?

- $\mathcal{R}(S)$ has $4g - 3$ topological components, which we label by integers $2 - 2g$ through $2g - 2$, corresponding to the Euler numbers of each component.
- $\mathcal{R}(S)$ is a real algebraic variety with two irreducible components.
- Each topological component contains a faithful representation. In particular:

  Theorem 19.6 (Deblois–Kent). *Faithful representations are dense in $\mathcal{R}(S)$.*
- Each algebraically irreducible component has the same dimension, and the singular locus is contained in the 0 component. All the components except for 0 are manifolds.
- The k-th component $\mathcal{R}_k$ is homeomorphic to $\mathcal{R}_{-k}$, but with the opposite orientation.

As for $\mathcal{H}(S)$, we have

$$\mathcal{H}(S) = \mathcal{H}_-(S) \sqcup \mathcal{H}_+(S),$$

where $\mathcal{H}_-(S)$ is the $2 - 2g$ component of $\mathcal{R}(S)$, and $\mathcal{H}_+(S)$ is the $2g - 2$ component.

20 2014-10-17: Hyperbolic surfaces, continued

Consider the Poincaré disk model, i.e., the open unit disk Δ with the metric

$$ds^2 = 4\frac{dx^2 + dy^2}{(1 - (dx^2 + dy^2))^2}.$$

This is isomorphic to the upper-half plane $\mathbb{H}$ with the metric $ds^2 = \frac{dx^2+dy^2}{y^2}$. Geodesic lines in Δ are straight lines through the center or circular arcs perpendicular to the boundary.

For $g \geq 2$, consider a regular $4g$-gon in Δ, whose boundary is given by $4g$ geodesics. We can continuously vary the polygon with a parameter $0 \leq \theta < \frac{(4g-2)\pi}{4g}$, the interior angle. This proves:

Lemma 20.1. *There exists a regular geodesic $4g$-gon in $\mathbb{H}$ with all angles $\frac{2\pi}{4g}$.*

Identifying opposite sides now yields a hyperbolic metric on a surface of genus g.

We can get more hyperbolic surfaces by gluing polygons together, using *pants decompositions*. Let P be a pair of pants with boundary components A, B, C. We want to build a hyperbolic metric on P such that the boundary components are locally geodesic. Moreover, we want to be able to specify lengths.

Every $S = S_{g,0}$ has a topological "pants decomposition". This would give us lots of hyperbolic metrics on S.

21 2014-10-20: Existence of hyperbolic structures

Let's continue with pants decompositions and Fenchel–Nielsen coordinates.

Lemma 21.1. *Any surface $S_{g,0}$ admits a pants decomposition. Every such decomposition has $2g-2$ pairs of pants. (Computing the number of distinct pants decompositions up to the action of* $\mathrm{Mod}(S)$ *is difficult and open.)*

Idea: Use the pants as building blocks for hyperbolic structures.

(1) Show how to build hyperbolic structures on S from hyperbolic structures on pants.

(2) Show every hyperbolic structure is like this.

We want a hyperbolic metric on a pair of pants P that has a locally geodesic boundary.

Definition 21.2. A submanifold N of a Riemannian manifold M is *totally geodesic* if every local geodesic tangent to N lies in N.

If boundary components of two pairs of pants have the same length, we can glue — in fact, we have a continuous family of gluings (given by the angle at which we glue the boundaries).

Between any two of the three boundary components, there is a unique minimal geodesic arc that intersects the boundaries at right angles. We can determine whether two gluings are the same by looking at the angle between the points on the shared boundary where these special arcs intersect the boundary on each pair of pants.

How do we build a hyperbolic metric with geodesic boundary? First, choose boundary lengths $2A, 2B, 2C \geq 0$. If we choose the three special arcs described above, then this will determine a hyperbolic metric. So, it remains to describe a right-angled hexagon R with ordered side lengths $A, *, B, *, C, *$. We can do this using bananas.

Exercise 21.3. This hexagon is uniquely determined (up to isometry) by A, B, C.

22 2014-10-22: Fenchel–Nielsen coordinates

Given two marked Riemann surfaces X, Y, how do we determine $d_{\mathcal{T}(S)}(X, Y)$, the Teichmüller distance?

Using Fenchel–Nielsen coordinates on $\mathcal{H}(S)$, a pants decomposition of (say) $S_{4,0}$ is given by 9 arcs. To give a hyperbolic structure, we give a length and a "gluing angle" for each arc, so there are 18 degrees of freedom for giving a hyperbolic metric.

Note that a 360° gluing gives an isometric surface. However, we still have a continuous family of hyperbolic metrics, looking something like $\mathbb{R}^9_+ \times (S^1)^9$.

It turns out that if we pass to the universal cover W of $\mathbb{R}^9_+ \times (S^1)^9$, then we get a homeomorphism $\mathcal{T}(S) \cong W$. The resulting coordinates $\mathbb{R}^9_+ \times \mathbb{R}^9$ are called *Fenchel–Nielsen coordinates*. In particular, we have maps

$$\mathcal{T}(S) \xrightarrow{\simeq} W \twoheadrightarrow \mathbb{R}^9_+ \times (S^1)^9 = \mathcal{T}(S)/\mathbb{Z}^9 \twoheadrightarrow \mathcal{M}(S) = \mathcal{T}(S)/\operatorname{Mod}(S).$$

Note that the $\mathbb{Z}^9$ is generated by Dehn twists around the pants-curves. Although $\operatorname{Mod}(S)$ is generated by Dehn twists, there are Dehn twists that don't come from the arcs in the pants decomposition. (In particular, $\operatorname{Mod}(S)$ has torsion.)

Even better:

Theorem 22.1 (Nielsen realization, Kerckhoff). *If $G \leq \operatorname{Mod}(S)$ is finite, then there exists $X \in \mathcal{T}(S)$ such that $g(X) = X$ for all $g \in G$.*

Why does this method get every hyperbolic surface homeomorphic to S? Because, given an essential simple closed curve γ in a hyperbolic surface X, γ is homotopic to a (local) geodesic.

23 2014-10-24

We've been talking about building hyperbolic structures directly from polygons and pants.

Since an essential simple closed curve has a (local) geodesic representative, we can pick a pants decomposition of a hyperbolic surface X, then isotope to a *geodesic* pants decomposition.

We got started by wanting to prove Teichmüller's theorem. We want to show existence of a Teichmüller mapping between any X and Y. The space $CQ'(X) \cong \mathbb{R}^{6g-6}$ is homeomorphic to $\mathcal{H}(S)$, the space of oriented hyperbolic structures on S, by uniformization:

$$u : CQ'(X) \xrightarrow{\simeq} \mathcal{H}(S).$$

We want to show that this is indeed surjective. Teichmüller's uniqueness theorem implies that u is injective. We'll show that $\mathcal{H}(S) \subset \mathbb{R}^{6g-6}$ is connected. Then we show that u is continuous and proper (inverse images of compact sets are compact). By Brouwer's invariance of domain, it follows that u is a homeomorphism.

Let's think about holomorphic representations. Consider $\operatorname{Hom}(\pi_1(S), \mathrm{PSL}_2 \mathbb{R})$ with the compact open topology (the topology of convergence on generators), where we think of $\pi_1(S)$ as a discrete group and $\mathrm{PSL}_2 \mathbb{R} = \mathrm{Isom}_+(\mathbb{H})$.

It would be nice if $\mathcal{H}(S) \subset \mathcal{R}(S)$. But we can always conjugate the image of a representation ρ by $A \in \mathrm{PSL}_2 \mathbb{R}$ to get a (possibly different) representation $\rho_A \in \mathcal{R}(S)$. But $\mathbb{H}/\rho(\pi_1(S))$ is isomorphic to $\mathbb{H}/\rho_A(\pi_1(S))$, so these representations give the same holomorphic structure. So we really need to work with $\mathcal{R}(S)/\mathrm{PSL}_2 \mathbb{R}$, where $\mathrm{PSL}_2 \mathbb{R}$ acts by conjugation.

What is $\mathcal{R}(S)$? We have

$$\pi_1(S) = \left\langle a_1, b_1, \ldots, a_g, b_g \;\middle|\; w_g := [a_1 b_1] \cdots [a_g, b_g] = 1 \right\rangle.$$

Then

$$\mathcal{R}(S) = \ker\left(\operatorname{Hom}(F_{2g}, \mathrm{PSL}_2 \mathbb{R}) = (\mathrm{PSL}_2 \mathbb{R})^{2g} \xrightarrow{w_g} \mathrm{PSL}_2 \mathbb{R}\right).$$

So $\dim_{\mathbb{R}} \mathcal{R}(S) = 6g-3$, and quotienting by $\mathrm{PSL}_2 \mathbb{R}$ reduces the dimension by 3 again (because $\mathrm{PSL}_2 \mathbb{R} \cong \mathbb{H} \times S^1$ has real dimension 3). This quotient map is bad, but one can check that it's fine over $\mathcal{H}(S)$.

Let's look at this naively. First, notice that elements $A \in \mathrm{PSL}_2 \mathbb{R}$ fall into three types:

elliptic $|\operatorname{tr}(A)| < 2$. Elliptic elements are conjugate in $\mathrm{PSL}_2 \mathbb{R}$ to "rotations", and are conjugate in $\mathrm{PSL}_2 \mathbb{C}$ to rotations around the origin.

parabolic $|\operatorname{tr}(A)| = 2$. Parabolic elements are conjugate in $\mathrm{PSL}_2 \mathbb{R}$ to $z \mapsto z + c$.

hyperbolic $|\operatorname{tr}(A)| > 2$. Hyperbolic elements are conjugate in $\mathrm{PSL}_2 \mathbb{R}$ to $z \mapsto \lambda z$.

Lemma 23.1. *If $\rho : \pi_1(S) \to \mathrm{PSL}_2 \mathbb{R}$ is discrete and faithful, then $\rho(g)$ is hyperbolic or the identity for all $g \in \pi_1(S)$.*

The hard part is showing that there are no parabolic elements since S is compact.

What are good coordinates for $\mathcal{H}(S)$? Take a discrete faithful representation $\rho : \pi_1(S) \to \mathrm{PSL}_2 \mathbb{R}$, and consider $\rho(a_1), \rho(b_1), \ldots, \rho(a_g), \rho(b_g)$.

24 2014-10-27

Let's show that $\Gamma = \rho(\pi_1(S))$ is "purely hyperbolic", as stated in the lemma last time. We want to normalize Γ. Let $\alpha_i = \rho(a_i)$ and $\beta_i = \rho(b_i)$, where we have presented $\pi_1(S)$ as

$$\pi_1(S) = \left\langle a_1, b_1, \ldots, a_g, b_g \;\middle|\; w_g := [a_1 b_1] \cdots [a_g, b_g] = 1 \right\rangle.$$

A hyperbolic element $A \in \mathrm{PSL}_2 \mathbb{R}$ is determined by its (ordered) fixed points and a real number λ.

(1) First step of normalization: Conjugate Γ in $\mathrm{PSL}_2 \mathbb{R}$ so that β_g fixes 0 and ∞.

(2) Second step: Conjugate so that α_g has fixed points $\lambda > 0$ and $-\frac{1}{\lambda}$.

Lemma 24.1. *After this normalization, Γ is now completely determined by the $6g-6$ real parameters $\varphi_i^-, \varphi_i^+, \lambda_{i,\alpha}, \psi_i^-, \psi_i^+, \lambda_{i,\beta}$, where $i \le 1 \le g-1$.*

So this allows us to explicitly consider $\mathcal{H}(S) \subset \mathbb{R}^{6g-6}$. This uses a lemma:

Lemma 24.2. *In a discrete subgroup of $\mathrm{PSL}_2\,\mathbb{R}$, two hyperbolic elements share a fixed point only if they share a power.*

25 2014-10-29: Beltrami differentials

We saw last time that $\mathcal{H}_+(S) \subset \mathbb{R}^{6g-6}$, where we think of $\mathcal{H}(S)$ as the space of normalized orientation-preserving holonomy representations of hyperbolic metrics on $S = S_{g,0}$.

Lemma 25.1. *None of the α_i, β_i for $1 \le i \le g-1$ stabilize ∞.*

Lemma 25.2. *The topology that $\mathcal{H}_+(S)$ inherits from $\mathbb{R}^{6g-6}$ is the same as the topology inherited from $\mathrm{Hom}(\pi_1(S), \mathrm{PSL}_2\,\mathbb{R})$.*

Consider the "Teichmüller existence" question: Why is $CQ'(X) \xrightarrow{u} \mathcal{H}_+(S)$ surjective?

The hardest part is continuity. Let's go back to deformations of Riemann surfaces. We've seen Teichmüller deformations: a quadratic differential gives a flat metric with horizontal and vertical foliations, and we stretch by $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$, giving a new conformal structure.

Eventually, one shows that $Q(X)$ is $T_X^* \mathcal{T}(S)$, but we now want to consider the *tangent* space. This is a space of *Beltrami differentials*.

While a quadratic differential looks like $\varphi(z)\,dz^2$, a Beltrami differential looks like $\mu(z)\frac{d\bar{z}}{dz}$, where $\mu(z)$ is a function in L^∞ with $\|\mu\|_\infty < 1$.

Consider a Riemann surface X with some fixed hyperbolic metric on it, and a uniformization $X = \mathbb{H}/\Gamma_X$. A Beltrami differential is an L^∞ function like this such that

$$\mu(z) = \mu(\gamma z)\frac{\overline{\gamma'(z)}}{\gamma'(z)},$$

where $\gamma \in \Gamma_X$.

Forget about X and Γ_X for a moment. On the unit disk Δ, given $\mu(z) \in L^\infty(\Delta)$ with $\|\mu\|_\infty < 1$, consider the differential equation $f_{\bar{z}} = \mu(z) f_z$. (If $\mu = 0$, then f is holomorphic.)

Theorem 25.3 (Morrey)**.** *Let $\mu \in B(\Delta) = (L^\infty(\Delta))'$. Then $f_{\bar{z}} = \mu(z) f_z$ has a homeomorphic solution $f : \Delta \to \Delta$ (i.e., f extends to a homeomorphism $\overline{\Delta} \to \overline{\Delta}$), unique after normalizing to fix $-1, i, 1$.*

26 2014-10-31: The Beltrami equation

Given a $\mathbb{C}$-valued L^∞ function μ with $\|\mu\|_\infty < 1$, we consider the *Beltrami equation*

$$f_{\bar{z}} = \mu(z) f_z.$$

There is a homeomorphic solution f_μ, and any homeomorphic solution is quasiconformal and will extend to a homeomorphism of $\overline{\Omega} \to \overline{\Omega}$. Moreover, this solution is unique after normalizing to fix 3 points on $\partial\Omega$, or to fix ∞ and 1 when $\Omega = \mathbb{C}$.

Theorem 26.1 (Morrey)**.** *The solution f_μ to the Beltrami equation depends continuously on μ.*

Theorem 26.2 (Ahlfors–Bers)**.** *In fact, f_μ depends real-analytically on μ.*

Instead, let's look for solutions $f^\mu : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ to the Beltrami equation that are 0 on $\hat{\mathbb{C}} \setminus \Omega$, no longer requiring the solution to be homeomorphic on the unit disk. This no longer sends the unit circle to itself; instead, f^μ sends the unit circle Δ to some Jordan curve, bounding a region Ω.

So we now have a solution f^μ defined on the whole Riemann sphere, and conformal on $\hat{\mathbb{C}} \setminus \Delta$.

Theorem 26.3 (Ahlfors–Bers)**.** *f^μ depends $\mathbb{C}$-analytically on μ.*

The Riemann mapping theorem implies existence of a homeomorphism $f : \Delta \to \Omega$ with $f_{\bar{z}} = 0$. Furthermore, we can normalize to prescribe 3 values on $\partial\Delta$. Morrey really proves that we can do this with any μ, and Ahlfors–Bers shows that there's a nice dependence on the parameters.

All of this goes through if we demand that

$$\mu(z) = \mu(\gamma z) \frac{\overline{\gamma'(z)}}{\gamma'(z)}$$

for all $\gamma \in \Gamma$, where Γ is a discrete subgroup of $\mathrm{PSL}_2 \mathbb{C}$. Define $B(\Gamma)$ to be the set of Beltrami differentials for Γ, i.e., all $\mu \in L^\infty$ such that $\|\mu\|_\infty < 1$ and μ satisfies the above equation.

What does this have to do with Teichmüller theory? Let $X = \mathbb{H}/\Gamma_X$ and $Y = \mathbb{H}/\Gamma_Y$ be hyperbolic surfaces, and let $F : X \to Y$ be a quasiconformal map. Then Γ_X and Γ_Y act on $\mathbb{H}$, and F lifts to a map $\tilde{F} : \mathbb{H} \to \mathbb{H}$ that is compatible with these actions.

27 2014-11-03

Marked quasiconformal homeomorphisms $X \to Y$ correspond to quasiconformal conjugacies $\Gamma_X \to \Gamma_Y$.

Fix a Fuchsian group $\Gamma_X \cong \pi_1(X) \cong \pi_1(S)$, where $X = \mathbb{H}/\Gamma_X$. Consider all quasiconformal conjugates $\Gamma_X^\mu := f_\mu \Gamma_X f_\mu^{-1}$ as μ ranges over

$$B(\Gamma_X) = \left\{ \mu \;\middle|\; \mu(z) = \mu(\gamma z) \frac{\overline{\gamma'(z)}}{\gamma'(z)} \right\}.$$

Since f_μ depends continuously on μ, we get a continuous map $B(\Gamma_X) \twoheadrightarrow \mathcal{H}_+(S)$. Using this gives us connectivity of $\mathcal{H}_+(S)$.

By Teichmüller's uniqueness theorem, the restriction of this map to $CQ(X)' \subset B(\Gamma_X)$ is injective, giving an injection $CQ(X)' \hookrightarrow \mathcal{H}_+(S)$. Invariance of domain implies that this map is open. However, this map is also closed, so it maps onto a connected component, hence is a homeomorphism.

28 2014-11-05: Teichmüller's existence theorem

Let $D(X)$ be the space of all quasiconformal deformations, $CQ^1(X)$ all Teichmüller deformations, $\mathcal{H}_+(S) \subset \mathbb{R}^{6g-6}$ all normalized holonomy representations of Fuchsian groups isomorphic to $\pi_1(S)$ and compatible with the orientation, and $u : CQ^1(X) \to \mathcal{H}_+(S)$ the map given by uniformization. Lift u to a map $\tilde{u} : D(X) \to \mathcal{H}_+(S)$.

We want u to be surjective, and this will give us Teichmüller existence. Every (marked) Y can thus be obtained from X by a Teichmüller deformation.

Picture: Given a quasiconformal map $f : X \to Y$, we get an equivariant quasiconformal map $\tilde{f} : \mathbb{H} \to \mathbb{H}$ and an equation

$$\tilde{f}_{\bar{z}} = \mu(z)\tilde{f}_z \tag{28.1}$$

with $\mu(z) \in (L^\infty(\mathbb{H}))^1$ such that

$$\mu(z) = \mu(\gamma z)\frac{\overline{\gamma'(z)}}{\gamma'(z)} \tag{28.2}$$

for all $\gamma \in \Gamma_X$. Conversely, given $\mu \in B(\Gamma) = \{\|\mu\|_\infty < 1,\ (28.2) \text{ holds}\}$, there is a solution $f : \mathbb{H} \to \mathbb{H}$ to (28.1) that is unique once normalized, and this descends to a quasiconformal map $f : X \to X_\mu$.

Note that $D(X) = B(\Gamma_X)$ is a ball about 0 in a vector space, hence is connected. Hence, $\mathcal{H}_+(S)$ is connected, i.e., for any marked Y, there is a smooth quasiconformal map $X \to Y$ in the correct homotopy class. Also, u is injective.

Now the hard part is over: Solutions f_μ depend continuously on μ. Representations depend continuously on f_μ (representations are changing by quasiconformal conjugacies), so the maps u are continuous.

By invariance of domain, $u : CQ^1(X) \hookrightarrow \mathcal{H}_+(S) \subset \mathbb{R}^{6g-6}$ is open. The last piece is to show that u is proper, hence closed (because the codomain is a nice Hausdorff space). Since $\mathcal{H}_+(S)$ is connected, we're done.

29 2014-11-07: Analytic point of view

The idea for properness of the uniformization map $u : CQ^1(X) \to \mathcal{H}_+(S)$ is that, given a compact set $\mathscr{C} \subset \mathcal{H}_+(S)$, the Teichmüller distance $d_{\mathcal{T}}(X, Y)$ is bounded over $\mathscr{C}$. This completes our sketch of the proof of Teichmüller's existence theorem.

Let's talk about an alternative definition of Teichmüller space. Let $\Gamma \subset \mathrm{PSL}_2 \mathbb{R}$ be a discrete subgroup. Take the space of Beltrami differentials

$$B_\Gamma(X) = \left\{ \mu(z) \in L^\infty(\mathbb{H}) \;\middle|\; \|\mu\|_\infty < 1,\ \mu(z) = \mu(\gamma z) \frac{\overline{\gamma'(z)}}{\gamma'(z)}\ \forall \gamma \in \Gamma \right\},$$

and let $f_\mu : \mathbb{H} \to \mathbb{H}$ be the normalized solution to the Beltrami equation $f_{\bar{z}} = \mu f_z$. Say that $\mu \sim_{\mathcal{T}} \nu$ (called *Teichmüller equivalence*) iff $\partial f_\mu = \partial f_\nu$ (i.e., f_μ and f_ν have the same values on the boundary). We now define Teichmüller space as $\mathcal{T}(\Gamma) := B(\Gamma)/\sim$.

Remark 29.1. If we lift a Dehn twist to the universal cover, we get a map with interesting behavior on the boundary. On the other hand, a quasiconformal map supported on a small piece of the space doesn't lift to interesting behavior on the boundary.

We could also use the normalized solution $f^\mu : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$, where we begin with μ on $\mathbb{H}$, then extend it by 0 to $\hat{\mathbb{C}}$. Then $f^\mu \sim f^\nu$ iff $f^\mu|_{\mathbb{R}} = f^\nu|_{\mathbb{R}}$. There's a theorem that this gives the same thing.

If $\mathbb{H}/\Gamma \cong S$, we get Teichmüller space $\mathcal{T}(S)$. But, we could also take $\Gamma = 1$ and get the *universal Teichmüller space* $\mathcal{T}(1) = \mathcal{T}(\mathbb{H}) = \mathcal{T}(\Delta)$. This has the property that $\mathcal{T}(\Gamma) \subset \mathcal{T}(\Delta)$ for all Γ.

The $\mathbb{C}$-vector space structure on $L^\infty(\mathbb{H})$ descends to a complex structure on $\mathcal{T}(\Gamma)$.

Conjugating Γ_X by f_μ gives a new Fuchsian group $\Gamma_\mu \subset \mathrm{PSL}_2 \mathbb{R}$. On the other hand, conjugating by f^μ gives a discrete subgroup $\Gamma'_\mu \subset \mathrm{PSL}_2 \mathbb{C}$, not necessarily lying in $\mathrm{PSL}_2 \mathbb{R}$. The subgroup Γ'_μ fixes some Jordan curve.

30 2014-11-10 [missing]

31 2014-11-12: Geometry of Teichmüller space

Let $S = S_{g,0}$, $g \geq 2$. Teichmüller space $\mathcal{T}(S)$ is a uniquely geodesic metric space homeomorphic to $\mathbb{R}^{6g-6}$.

What does $\mathcal{T}(S)$ look like geometrically? We can ask global questions, like:

- Is $\mathcal{T}(S)$ hyperbolic in any sense?
- What is $\mathrm{Isom}(\mathcal{T}(S))$? It's the mapping class group $\mathrm{Mod}(S)$ acting by changing markings. (Royden)
- Given X and Y, can we "find" the geodesic between them?
- What does it mean to leave a compact set?
- Is $\mathcal{T}(S)$ isometric (or quasi-isometric) to something we know?

Also, local questions:

- Is $\mathcal{T}(S)$ homogeneous, i.e., does it look the same at any two points X, Y? No, it's highly non-homogeneous: If X and Y have small isometric neighborhoods, then $gX = Y$ for some $g \in \mathrm{Mod}(S)$.

- Are metric balls convex?
- Is the metric Riemannian? No, but it is *Finsler*, i.e., there's a Minkowski norm on each tangent space.

There's an open problem related to convexity of metric balls.

Definition 31.1. The *convex hull* of a subset is the smallest convex set containing it.

Open problem: Can the convex hull of 3 points be all of $\mathcal{T}(S)$?

The local behavior is very hard to study, and we'll mostly ignore it. Let's try to describe the metric globally. Think of $\mathcal{T}(S)$ as a space of hyperbolic surfaces with geodesics given by analysis (Teichmüller deformations and flat metrics).

Theme: Use simple closed curves on the surface (e.g. geodesics on hyperbolic surface). Consider the moduli space $\mathcal{M}(S) = \mathcal{T}(S)/\operatorname{Mod}(S)$. Let

$$\mathcal{M}_\varepsilon(S) = \left\{ X \in \mathcal{M}(S) \;\middle|\; \inf \ell(\gamma) \geq \varepsilon \right\},$$

where γ ranges over essential simple closed curves on S. (A simple closed curve is *essential* if it doesn't contract to a point or a puncture.) We call this the "ε-thick part of $\mathcal{M}(S)$".

Theorem 31.2 (Mumford)**.** *$\mathcal{M}_\varepsilon(S)$ is compact.*

So, leaving every compact set in $\mathcal{M}(S)$ means some curve is getting short.

32 2014-11-14: Grothendieck's perspective

What is a deformation of Riemann surfaces? Think of it as a family of varieties (given, in a simple case, by moving the coefficients of a defining equation around).

We can also add more constraints, such as the varieties all being homeomorphic, all algebraic curves, etc. What is a family? A family of Riemann surfaces is a holomorphic fibration over a $\mathbb{C}$-manifold (or variety) with curves as fibers.

A moduli space is a *classifying space for families.* Studying families of curves is equivalent to studying maps into the moduli space.

Grothendieck's "Lego" perspective is to piece together moduli problems from pieces of smaller complexity curves. There's some even more universal moduli space coming from piecing together data from all Riemann surfaces at once.

In a Riemann surface, let a curve γ get short. Actually, contract its length all the way to zero, and make γ part of a pants decomposition. This attaches a copy of $\mathcal{T}(S_{g-1,n+2})$ at the boundary of $\mathcal{T}(S_{g,n})$. This *doesn't* compactify $\mathcal{T}(S)$, but it gives a "bordification".

Grothendieck uses this perspective in "Esquisse d'un Programme" to approach the inverse Galois problem using what he referred to as *anabelian geometry*.

32.1 Short curves

The boundary components mentioned above intersect in a complicated way that's tracked by a simplicial complex.

If we have a loop γ in X that is short, i.e., $\ell_X(\gamma)$ is small, then γ has a *wide collar* in X.

Lemma 32.1 (Keen–Halbern collar lemma)**.** *For any γ, there is a collar of width*

$$\operatorname{arcsinh}\left(\frac{1}{\sinh(\frac{\ell}{2})}\right).$$

33 2014-11-17: Collars

Lemma 33.1 (Collar lemma, Keen–Halpern)**.** *In a hyperbolic surface, short geodesics have big collars. (Long geodesics have "shallow" collars.)*

By a "collar", we mean a tubular neighborhood of the curve that can be deformation-retracted to the curve.

In both cases, we can bounds the "width" or "depth" of the collar in terms of length, as stated last time.

Corollary 33.2. *If two embedded geodesic loops γ and δ intersect and $\gamma \neq \delta$, then they can't both be short.*

There's no reason we'll have a short curve. On the other hand, there are always "moderate-length" curves:

Theorem 33.3 (Bers)**.** *Let $n \in \mathbb{N}$. There there is a constant $B = B(n)$ such that, if X is a compact surface of genus $g \leq n$, then there is a geodesic loop γ in X such that $\ell_X(\gamma) \leq B$.*

34 2014-12-01: Thin-parts and the curve complex

Given a closed surface S (possibly with punctures), there is a simplicial complex $\mathscr{C}(S)$ defined as follows: The 0-skeleton is

$$\mathscr{C}^{(0)}(S) = \{\text{isotopy classes of essential simple closed curves in } S\},$$

and the k-skeleton $\mathscr{C}^{(k)}(S)$ is defined by declaring that $\gamma_0, \ldots, \gamma_k$ span a k-simplex if they can be realized disjointly.

Remark 34.1. $\mathscr{C}(S)$ is a *flag complex*, i.e., $\mathscr{C}(S)$ is determined by its 1-skeleton $\mathscr{C}^{(1)}(S)$.

Remark 34.2. $\mathscr{C}(S)$ is finite-dimensional. Indeed, once we lay down enough curves to give a pants decomposition, any essential simple closed curve will either be isotopic to or have nonempty intersection with one of the curves in the pants decomposition. In particular, the dimension of $\mathscr{C}(S)$ is one less than the number of curves in a pants decomposition.

Remark 34.3. $\mathscr{C}(S)$ is not locally finite; Dehn twists can produce infinite families of isotopy classes of curves that are disjoint from a given curve. Nor is $\mathscr{C}(S)$ locally compact.

Remark 34.4. The mapping class group $\mathrm{Mod}(S)$ acts simplicially on $\mathscr{C}(S)$, giving a quotient $\mathscr{C}(S)/\mathrm{Mod}(S)$, which is compact because, up to $\mathrm{Mod}(S)$, there are only finitely many pants decompositions.

Theorem 34.5 (Serre–Ivanov). *The subgroup* $\Gamma_3(S) := \ker(\mathrm{Mod}(S) \to \mathrm{Sp}_{2g}(\mathbb{Z}/3\mathbb{Z}))$ *is* pure, *i.e., if* $\gamma \in \Gamma_3$ *fixes a simplex of* $\mathscr{C}(S)$*, then* γ *fixes the simplex pointwise.*

So far, we're considering $\mathscr{C}(S)$ as a complex with the *weak topology.* This leads to theorems such as:

Theorem 34.6 (Harer). *With the weak topology,* $\mathscr{C}(S)$ *is homotopy equivalent to a nonempty wedge of spheres, all of the same dimension.*

This tells us things about $\mathrm{Mod}(S)$, $\mathcal{M}(S)$, etc. For example, the stabilizers in $\mathrm{Mod}(S)$ are mapping class groups of subsurfaces.

There's also a metric topology on $\mathscr{C}(S)$. Any simplicial complex K has an induced path metric: distances are given by the shortest length of any piecewise-linear path in K. (In other words, we consider each simplex to be a regular Euclidean simplex, and glue together these metric spaces.)

Theorem 34.7. *This metric topology on* $\mathscr{C}(S)$ *is weakly equivalent to the weak topology.*

Theorem 34.8 (Kobayashi). *With the metric topology,* $\mathscr{C}(S)$ *is an infinite-diameter metric space.*