Metrics and Coördinates on Teichmüller Space

Feng Zhu

July 28, 2017

These are summary notes for a minicourse given at the University of Michigan in May 2017.

0 Teichmüller space; a brief overview

We very briefly recall what Teichmüller space is.

Let Σ_g be a hyperbolizable surface of genus g, i.e. an orientable 2-(real-)dimensional manifold with $\chi(\Sigma_g) < 0$, or equivalently $g \geq 2$.

Except in Section 5, Σ_g will be closed, i.e. compact and without boundary.

Definition 0.1. The genus-g Teichmüller space, denoted $\mathcal{T}_g = \mathcal{T}(\Sigma_g)$, is the set of marked hyperbolic[1] metrics on Σ_g up to isotopy.

By uniformization, this is equivalent to the set of marked conformal structures or marked Riemann surface structures on Σ_g up to isotopy.

A point in $\mathcal{T}_g$ can be specified as a(n equivalence class of) pair(s) $[(S, h)]$ where S is a hyperbolic (or Riemann, depending on the point of view being taken) surface, and $h : \Sigma_g \to S$ is an orientation-preserving homeomorphism.

The equivalence relation is then defined as follows: $(S, h) \sim (S', h')$ if there exists an isometry (or biholomorphism, if we're thinking of the conformal point of view) $j : S \to S'$ s.t. the diagram

$$\begin{array}{ccc} \Sigma_g & \xrightarrow{h} & S \\ & \searrow^{h'} & \downarrow j \\ & & S' \end{array}$$

commutes up to isotopy.

More evocatively: think of Σ_g as the "naked" topological surface, S as the hyperbolic (or conformal) "clothing", and the marking h as "instructions for how to wear the clothing." The same (isometric / biholomorphic) clothing, with instructions which differ only up to a small jiggle (isotopy), is considered equivalent.

Aside: the set of (unmarked) hyperbolic metrics or conformal structures on Σ_g up to isotopy is **moduli space** $\mathcal{M}_g$, which is the quotient of $\mathcal{T}_g$ by the **mapping class group** Mod_g.

(I am not defining these here; if you would like to read more, see e.g. Benson Farb and Dan Margalit's *Primer on Mapping Class Groups*.)

The mapping class group is big and complicated and all sorts of fun, and its rich structure is reflected in how much trickier $\mathcal{M}_g$ is to study compared to $\mathcal{T}_g$. This is one reason Teichmüller space is studied even though, in some senses, moduli space is a more natural object.

To make sense of $\mathcal{T}_g$ as a space and not just a set, we need to put a topology (and maybe also geometry, eventually) on it. We would like to do so, if possible, in a way which "naturally" reflects the structure of $\mathcal{T}_g$ as a set of hyperbolic metrics (or conformal structures), instead of being completely arbitrary.

There are a number of ways of doing so: e.g. via any of a number of coordinate systems or metrics which we will describe below[2], or using the algebraic topology (see Section 5 below.)

[1] i.e. Riemannian with constant sectional curvature -1

[2] The former give $\mathcal{T}_g$ the topology of a ball; the latter induce metric topologies. All of these, as well as the algebraic topology, will turn out to be the same topology.

A brief outline of the contents:

- we first introduce two more classical constructions: Fenchel-Nielsen coordinates, which use the hyperbolic point of view, and the Teichmüller metric, which uses the conformal point of view;
- then we take a look at the Weil-Petersson metric—which uses both points of view, but is somehow much more mysterious—and its various interpretations;
- then we introduce the Thurston metric, which is analogous to the Teichmüller metric but uses the hyperbolic point of view, and at shearing coordinates, which came from Thurston's study of this metric.
- We then briefly remark on the various metrics on Teichmüller space which come from its complex manifold structure,
- and touch on Penner's notion of λ-lengths, which gives coordinates for Teichmüller spaces of *punctured* hyperbolizable surfaces; these coordinates have many nice properties and deep connections to many other areas of mathematics.
- Finally, we end with a description of some applications of these various coordinate systems and metrics.

We caution the reader that we will be (in these notes, as we were in the minicourse) somewhat vague throughout, and only seek to capture and transmit broad-stroke ideas of things; references to places where more careful / detailed accounts can be found will be included throughout.

1 Fenchel-Nielsen coordinates and the Teichmüller metric

1.1 Fenchel-Nielsen coordinates

Theorem 1.1 (Fenchel-Nielsen)**.** *$\mathcal{T}_g$ is homeomorphic to* $\mathbb{R}^{6g-6} = \mathbb{R}^{3|\chi(\Sigma_g)|}$.

Sketch of proof. We construct an explicit homeomorphism to $\mathbb{R}^{3g-3} \times \mathbb{R}_+^{3g-3}$ **using hyperbolic geometry**:

- Decompose surface into hexagons:
  - Cut Σ_g into $2g-2$ pairs of pants (three-holed spheres.)
  - Cut each pair of pants into two hexagons.

    If the pair of pants is equipped with a hyperbolic metric s.t. the boundary is geodesic, cut along the common perpendiculars[3] between each pair of boundary components (the "seams" of the pants. This produces two all-right-angle hexagons.
- An all-right-angle hyperbolic hexagon is uniquely determined (up to isometry) by specifying every other side-length, and we may specify these three lengths arbitrarily.

  In particular, we may specify the boundary lengths of each pairs of pants arbitrarily, and specifying them uniquely determines the hyperbolic structure on the pair of pants.

  There are $3g-3$ curves in a pants decomposition[4] (each of the $2g-2$ pairs of pants has 3 boundary components, but these $6g-6$ curves are identified in pairs), and by the above we obtain $3g-3$ "length" coordinates in $\mathbb{R}_+$.
- We may glue together hyperbolic pairs of pants as long as the boundary lengths agree (i.e. given a surface Σ, possibly with boundary, that decomposes into pairs of pants, and a choice of hyperbolic metrics on each pair of pants s.t. corresponding boundary lengths agree, there is a [unique] hyperbolic metric on Σ which restricts to the specified hyperbolic metric on each pair of pants.)

[3] This is an essential fact from hyperbolic geometry that we are using here: given two disjoint geodesics in $\mathbb{H}^2$ (which do not share boundary points either), there exists a unique common perpendicular between them.

[4] i.e. the set of curves one cuts Σ_g along to end up with $2g-2$ pairs of pants.

There is an additional degree of freedom which comes in when we glue: we can "twist" one curve relative to the other so that the seams are offset.

(Changing the twist coordinate preserves the hyperbolic metric in a neighborhood of the curve around which the twist occurs, but changes e.g. the lengths of geodesics transversely intersecting said curve / neighborhood.) One can verify that negative and positive twists lead to different metrics, as does twisting by one whole revolution (so that the seams are once again aligned) as opposed to not twisting at all.

Hence we get $3g-3$ "twist" coordinates in $\mathbb{R}$.

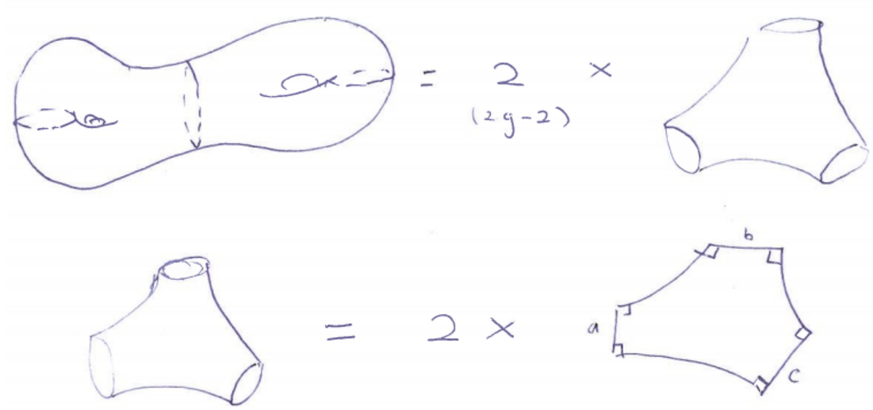

Figure 1: Cutting a (genus-2) surface into pairs of pants, and then all-right heaxgons.

We now **claim** that (fixing a pants decomposition) the lengths and twists give a homeomorphism from $\mathcal{T}_g$ to $\mathbb{R}_+^{3g-3} \times \mathbb{R}^{3g-3} \cong \mathbb{R}^{6g-6}$.

To verify this we

- build an explicit inverse (given a tuple of lengths and twists, we build hyperbolic pairs of pants with the specified boundary lengths, then glue them up with the specified twists.)
- check that the map and its inverse are both continuous, with respect to any of the topologies on $\mathcal{T}_g$ described above.

To do all of this rigorously (especially the bit involving the twist; also checking compatibility between the various topologies) takes a bit of work, but is not terrifically enlightening. For details, see e.g. Farb-Margalit's *Primer* (§10.6), or Yoichi Imayoshi and Masahiko Taniguchi's *Introduction to Teichmüller spaces* (§3.2.) □

Remark 1.2. The $6g-6$ numbers given by the lengths and twists are called Fenchel-Nielsen coordinates.

Changing the pants decomposition changes the homeomorphism / coordinates, but induces the same topology.

Fenchel-Nielsen coordinates are not canonical (since they depend on the choice of pants decomposition), although see Theorem 2.1.

1.2 Teichmüller metric

We can also give $\mathcal{T}_g$ a "natural" geometry by **defining a metric which measures distortion between conformal structures**.

Definition 1.3 (Quasiconformal maps)**.** Given Ω, Ω' open sets in $\mathbb{C}$, an orientation-preserving homeomorphism $f : \Omega \to \Omega'$ is k**-quasiconformal** if f is absolutely continuous on lines[5] and $K_f = \frac{1+|\mu_f|}{1-|\mu_f|} \leq k$ a.e., where $\mu_f := \frac{f_{\bar{z}}}{f_z}$.

[5]A regularity condition which implies in particular that the derivatives f_z and $f_{\bar{z}}$ exist a.e.—this is the key thing that is really needed here, so that the next part of the definition makes sense.

μ_f is called the **Beltrami differential** associated to f, and it is really a $(-1,1)$-form.

K_f is called the **(complex) dilatation** of f. Both of these quantities measure the distortion of the conformal structure, in the following precise sense:

Proposition 1.4. *If f is differentiable at p, $df(T_p^1\Omega)$ is an ellipse with aspect ratio $K_f(p) = \frac{|f_z|+|f_{\bar{z}}|}{|f_z|-|f_{\bar{z}}|}(p)$, and the major axis of the ellipse makes an angle of $\frac{1}{2}\arg\mu_f(p)$ with the real axis.*

Given S_1, S_2 homeomorphic hyperbolic surfaces, $f : S_1 \to S_2$ is called k-quasiconformal iff it lifts to a k-quasiconformal map $\tilde{f} : \mathbb{H}^2 \to \mathbb{H}^2$ (where $\mathbb{H}^2$ is viewed, via e.g. the upper half-plane model, as an open set in $\mathbb{C}$.) This is equivalent to requiring that the restriction of f to (conformal) coordinate charts be k-quasiconformal

Definition 1.5 (Teichmüller metric)**.** Given $p_1 = [(S_1, h_1)], p_2 = [(S_2, h_2)] \in \mathcal{T}_g$, define

$$d_{Teich}(p_1, p_2) := \inf\{\log k : \exists f : S_1 \to S_2 \; k\text{-qc}, f \circ h_1 \simeq h_2\}$$

We remark that this is well-defined (under the equivalence relation defining Teichmüller space) since pre- or post-composing a k-quasiconformal map with a conformal map yields a k-quasiconformal map.

Theorem 1.6 (Teichmüller existence and uniqueness)**.** *Given $p_1 = [(S_1, h_1)], p_2 = [(S_2, h_2)] \in \mathcal{T}_g$, there* **exists** *a* **unique** *$h : S_1 \to S_2$ s.t. $h \simeq h_2 \circ h_1^{-1}$ and h is k-quasiconformal where $k = d_{Teich}(p_1, p_2)$.*

Corollary 1.7. *d_{Teich} is a metric.*

Proof. This is really a consequence of the properties of quasiconformal maps, together with Teichmüller existence:

For symmetry, we observe that a homeomorphism h is k-quasiconformal iff h^{-1} is k-quasiconformal.

For the triangle inequality, we observe that if h_1 is k_1-quasiconformal and h_2 is k_2-quasiconformal, then $h_2 \circ h_1$ is [at most] k_2k_1-quasiconformal.

It is obvious that $d(p_1, p_2) \geq 0$, since the quasiconformality constant is never < 1, and also that $d(p, p) = 0$; to show that we have equality only if $p_1 = p_2$, we use Teichmüller existence to say that this implies the existence of a 1-quasiconformal map $h : S_1 \to S_2$—but then 1-quasiconformal maps are conformal. □

Idea of proof for Teichmüller theorems. In genus 1, the "best" / "least-distortion" map φ between conformal structures are **affine** (a result due to Grötzsch), i.e. there exist pairs of transverse measured foliations on the source and target surfaces s.t. φ takes one pair to the other, stretches along the leaves of one, and contracts along the leaves of the other.

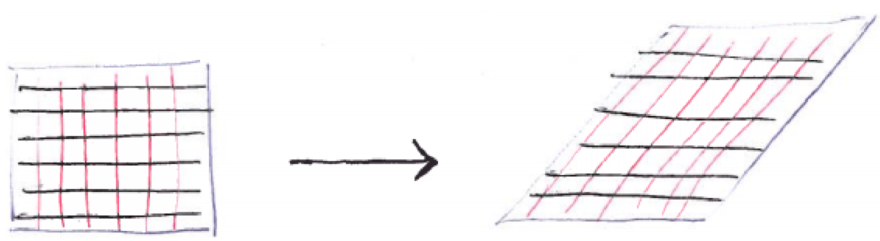

Figure 2: An affine map on a torus

In higher genus, they are almost affine: now we have pairs of *singular* transverse measured foliations, which look just like regular transverse foliations away from some finite number of singularities

x and we can similarly define maps ("Teichmüller mappings") which take one pair to the other, send singularities of one to singularities of the other, and away from the singularities stretches along the leaves of one set of foliations and contracts alongs the leaves of the other.

We remark, in passing, that these foliations can be encoded using **holomorphic quadratic differentials** (i.e. $(2,0)$-forms) q—the zeroes of q are where the singularities are; the leaves of one foliation are paths tangent to vectors which evaluate to positive reals under q, the leaves of the other are paths tangent to vectors which evaluate to negative reals.

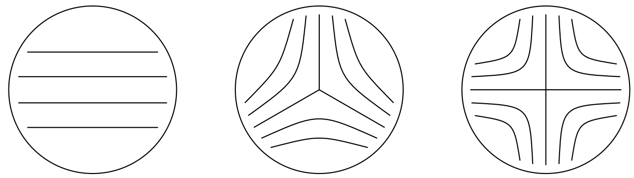

Figure 3: Local pictures (left) away from a singular point; (center) at a 3-prong singularity; (right) at a 4-prong singularity.

Then, **for Teichmüller uniqueness**, we show that any other quasiconformal map in the homotopy class has a larger quasiconformality constant.

For Teichmüller existence, we build a surjective map Θ from the vector space of holomorphic quadratic differentials on S_1 (the source surface) to $\mathcal{T}_g$—the Teichmüller mapping along the set of foliations encoded by the (a) differential corresponding to S_2 is the map we seek. □

For further details we refer the reader to Imayoshi-Taniguchi (§4 and §5.)

Some properties of d_{Teich}:

- Finsler (comes from / induces a continuously varying family of norms on the tangent spaces), not Riemannian (but not inner products)
- complete (geodesics exist for all time, in both directions)
- geodesic; geodesics are Teichmüller lines
- not negatively-curved: more specifically,
  - (Masur) not negatively-curved in the sense of Busemann (i.e. there exist geodesic rays starting from the same point which do not diverge exponentially)
  - (Masur-Wolf) not Gromov-hyperbolic (i.e. geodesic triangles are not "thin" at any scale)

  However (Masur-Minsky) $\mathcal{T}_g$ with the Teichmüller metric exhibits a more subtle / complicated variant of negative / non-positive curvature, hierarchical hyperbolicity.
- Mod_g-invariant; in fact (Royden) $\text{Isom}^+(\mathcal{T}_g, d_{Teich}) = \text{Mod}_g$ (except when $g = 2$, where the isometry group is an index-2 extension of Mod_g by the hyperelliptic involution.)

2 The Weil-Petersson metric

2.1 A (quick and rough) definition; Properties

Can we build a Riemannian metric? Yes we can.

1. Identify the (co)tangent spaces with vector spaces of complex analytic objects:

$$T_{[(S,h)]}\mathcal{T}(\Sigma_g) = BD_1(S)/BD_0(S)$$
$$T^*_{[(S,h)]}\mathcal{T}(\Sigma_g) = QD_1(S)$$

   where $BD_1(S)$ denotes the space of (essentially bounded) Beltrami coefficients on S, $BD_0(S)$ denotes the subspace of homotopically-trivial ones, and $QD_1(S)$ denotes the space of (essentially bounded) holomorphic quadratic differentials on S.

   The identification proceeds e.g. via the map

$$BD_1(S) \to \{\text{quasiconformal homeos from } S\} \stackrel{\Phi}{\to} \mathcal{T}(\Sigma_g)$$

where the first part is given by the measurable Riemann mapping theorem, and the second part is given by $\varphi \mapsto \varphi(S)$. Then we have

$$BD_1(S) \to = T_0\{\text{quasiconformal homeos from } S\} \overset{d\Phi}{\to} T_{[(S,h)]}\mathcal{T}(\Sigma_g)$$

and the kernel of $d\varphi$ is given precisely by $BD_0(S)$.

Since (from the above) both Beltrami differentials and holomorphic quadratic differentials capture, in various ways, [infinitesimal] distortion of the conformal structure, this identification should not be too surprising.

2. Define the metric on tangent spaces as follows: on $T_{[(S,h)]}\mathcal{T}(\Sigma_g)$,

$$g_{WP}(\mu_1, \mu_2) = \int_S \mu_1 \bar{\mu_2}\, ds^2$$

where ds^2 denotes the hyperbolic area element on S—this is a L^2 inner product on the Hilbert space $BD_1(S)/BD_0(S)$.

(We can also define the cometric on the cotangent spaces by

$$g^*_{WP}(q_1, q_2) = \int_S q_1 \bar{q_2}\, ds^{-2},$$

which more closely reflects certain inspirations from modular forms.)

The real part of this Hermitian metric is then the Riemannian metric we seek.

Some properties of d_{WP}: (compare these to the properties of d_{Teich} above)

- Riemannian
- **not** complete (Chu-Wolpert) BUT geodesically convex (which is good enough for some arguments)
  The metric completion is augmented Teichmüller space $\bar{\mathcal{T}}(\Sigma_g)$, which is CAT(0).
- geodesics are ... mysterious
- (Wolpert) has negative sectional and Ricci curvatures BUT curvatures not bounded away from $-\infty$ or 0.
- Mod_g-invariant: (Masur-Wolf) $\text{Isom}^+(\mathcal{T}_g, d_{WP}) = \text{Isom}^+(\mathcal{T}_g, d_{Teich})$.
- Kähler[6]

The Weil-Petersson metric, or more precisely its associated symplectic form ω_{WP} (= imaginary part of g_{WP}), can be used to show that the construction leading to Fenchel-Nielsen coordinates is, in some sense, canonical:

Theorem 2.1 (Wolpert's magic formula)**.** *Take* **any** *pants decomposition of Σ_g. Then*

$$\omega_{WP} = \sum_{i=1}^{3g-3} d\ell_i \wedge d\tau_i$$

where ℓ_i and τ_i are the length and twist functions.

A key step in the proof of this is the following statement:

[6] Claimed by Weil in 1958, proven by Ahlfors in 1961.

Theorem 2.2 (Wolpert's twist-length duality)**.** *Let C be a simple closed curve, ℓ_C be the corresponding length function on Teichmüller space, and t_C be the vector field associated to the path in Teichmüller space given by twisting along C. Then*

$$\omega_{WP}(t_C, \cdot) = -d\ell_C.$$

Wolpert also used this statement, with more work, to produce things such as sine-length and cosine formulas and curvature estimates.

The general theme here seems to be that "the hyperbolic geometry on the surface is reflected in the symplectic geometry of Teichmüller space (with the Weil-Petersson metric.)"

For further details, see e.g. §8 of Imayoshi-Taniguchi or §5 of Harer.

The original formulation of the Weil-Petersson metric, although it makes nodding use of the hyperbolic metric, is arguably essentially analytic / conformal in nature. There are other ways of approaching it which more explicitly / exclusively start from the hyperbolic point of view. Below we present two of these.

We remark that, historically, these were originally conceived of as alternative ways of putting metrics on Teichmüller space, and then realized to be equivalent to the Weil-Petersson metric.

2.2 Reinterpretation as "length of a random geodesic"

Motivating idea: Given a homotopy class of closed curve $C \subset \Sigma_g$ (i.e. element of $\pi_1\Sigma_g$) and S a genus-g hyperbolic surface, let $\ell_C(S)$ be the length of (the unique geodesic representative of) C, according to the hyperbolic metric S.

For any given C, we think of ℓ_C as a function on Teichmüller space, and now we look at its behavior along certain special paths in Teichmüller space, called **twist paths**. A twist path is given by picking out some (multi)curve on a hyperbolic surface, and "twisting" the hyperbolic surface along that (multi)curve, i.e. in the language of Fenchel-Nielsen coordinates: increasing the twist coordinate along the (multi)curve.

Fact 2.3. $\ell_C : \mathcal{T}(\Sigma_g) \to \mathbb{R}_{>0}$ is strictly convex along twist paths, and bounded from below.

Fact 2.4 (Thurston's Earthquake Theorem, super-vague version)**.** Any two points in Teichmüller space are connected by a suitably generalized twist path.

And so, the idea goes, the Hessian of ℓ_C at a minimum should give a positive-definite quadratic form, i.e. a Riemannian metric.

Now comes the **problem**: which C should we pick?

Thurston offered an answer, which in essence says we should pick C to be a limit of curves "distributed uniformly at random" on our [basepoint] hyperbolic surface. Below we will make this slightly precise, using the framework and language of geodesic currents as introduced by Bonahon.

Definition 2.5. A **geodesic current** on Σ_g is a $\pi_1(\Sigma_g)$-invariant measure on the space $(\partial_\infty \pi_1(\Sigma_g))^{(2)}$ of unordered pairs of distinct points on the Gromov boundary (given a hyperbolic metric on Σ_g, this can be identified with the space of unoriented geodesics on Σ_g.)

Write $\mathcal{C}(\Sigma_g)$ to denote the space of geodesic currents on Σ_g.

We remark that the use of the word "space" implies a topology on these sets, which we are not specifying here.

In order to simplify the language below we fix an arbitrary hyperbolic metric on Σ_g (think of it as choosing a "basepoint" in Teichmüller space), so that we can talk about closed geodesics instead of homotopy classes of closed curves, and identify $\pi_1(\Sigma_g)$ with the set of closed geodesics on Σ_g.

The space of closed geodesics on Σ_g embeds into $\mathcal{C}(\Sigma_g)$, by sending a curve C to the atomic measure δ_C. The image is in fact dense in $\mathcal{C}(\Sigma_g)$. Perhaps slightly more surprising is the following:

Theorem 2.6 (Bonahon)**.** *Teichmüller space $\mathcal{T}(\Sigma_g)$ embeds into $\mathcal{C}(\Sigma_g)$, by sending a hyperbolic metric S to a unique associated measure L_S, known as the Liouville measure of S.*

With these identifications, we have:

Theorem 2.7 (Bonahon)**.** *The intersection form $i : \pi_1(\Sigma_g) \times \pi_1(\Sigma_g) \to \mathbb{R}_{>0}$ extends continuously to a symmetric bilinear form on $\mathcal{C}(\Sigma_g) \times \mathcal{C}(\Sigma_g)$.*

Given $C \in \pi_1(\Sigma_g)$ and $S \in \mathcal{T}(\Sigma_g)$, $i(C, L_S) = \ell_C(S)$.

Bonahon also shows in what sense the Liouville current may be thought of as "a random geodesic": given $v \in T^1\Sigma_g$, let α_t be a closed curve constructed as follows: take the geodesic arc of length t starting at v (i.e. starting at the footpoint of v and tangent to v at that point), join the endpoints of this arc by an arbitrary arc of length bounded above by some constant depending only on the choice of hyperbolic metric S on Σ_g), and take the geodesic representative of the resulting closed curve. One could see why the result of applying this construction to a "randomly-chosen" v might be thought of as "a random geodesic". In particular, if there is something we can say about α_t for almost every choice of v, then we might reasonably attach those things to "the random geodesic on Σ_g [with our choice of hyperbolic metric]."

Using the Birkhoff ergodic theorem and the ergodicity of the geodesic flow (among other things), one can show that

$$\lim_{t\to\infty} \frac{\alpha_t}{\ell_S(\alpha_t)} = \left(\pi\sqrt{|\chi(\Sigma_g)|}\right)^{-1} L_S$$

(in particular, the limit exists) for a.e. $v \in T_1\Sigma_g$. Here everything is done in the space of currents (so e.g. α_t is identified with δ_{α_t}.) Thus the Liouville current, up to a normalization constant depending only on the topology of Σ_g, can be thought of as a [normalized] "random geodesic".

Then we have

Theorem 2.8 (Thurston, Bonahon)**.** *$i : \mathcal{T}(\Sigma_g) \times \mathcal{T}(\Sigma_g) \to \mathbb{R}_{>0}$ (given by $(S_1, S_2) \mapsto i(L_{S_1}, L_{S_2})$) is a metric on $\mathcal{T}(\Sigma_g)$.*

Theorem 2.9 (Wolpert, Bonahon)**.** *This metric is a scalar multiple of the Weil-Petersson metric. More precisely,*

$$\frac{\partial^2}{\partial t \partial u} i(L_S, L_{S_{tu}})\big|_{t=u=0} = \frac{\pi}{3} d_{WP}\left(\frac{\partial}{\partial t} S_{tu}, \frac{\partial}{\partial u} S_{tu}\right).$$

for $(t,u) \mapsto S_{tu}$ any smooth 2-parameter family in $\mathcal{T}(\Sigma_g)$, with $S := S_{00}$.

For details, we refer the reader to Francis Bonahon's paper, "The geometry of Teichmüller space via geodesic currents".

2.3 Reinterpretation using thermodynamical formalism

(Note: this was not talked about during the minicourse.)

Given a hyperbolic surface S homeomorphic to Σ_g, we can associate to it the geodesic flow on the unit tangent bundle T^1S, and then associate to this flow a Hölder reparametrization function on T^1S.

From dynamics we take the notion of **pressure**, which is supposed to be a generalisation of entropy in the presence of some potential function[7] (here, associated to the reparametrization.)

Pressure is an analytic function on the space of α-Hölder functions. The Hessian of the pressure, restricted to the subspace of pressure-zero functions, can be shown to be non-degenerate (by length rigidity e.g. in the form of the $9g-9$ theorem, since along degenerate directions, no lengths change.)

We can take take our associated Hölder functions to be pressure-zero, and so this whole machine produces a metric on Teichmüller space.

Theorem 2.10 (McMullen)**.** *This metric is equal to the "random geodesic" metric described above, and hence a scalar multiple of the Weil-Petersson metric.*

For further details we refer the reader to the survey article of Martin Bridgeman, Richard Canary and Andrés Sambarino entitled "An introduction to pressure metrics on higher Teichmüller spaces".

We remark also that McMullen uses this re-interpretation to relate d_{WP} to dynamical quantities such as the Hausdorff dimension of the limit set, and to explore parallel / analogous results in complex dynamics: for details there, see his paper, "Thermodynamics, dimension and the Weil-Petersson metric".

[7] The terminology is from statistical mechanics, and some have argued that "pressure" should really be called "free energy".

3 The Thurston metric and shearing coordinates

3.1 The Thurston metric

The idea here is very similar to that leading to the Teichmüller metric. The difference is this: what if we defined a metric which measured hyperbolic instead of conformal distortion?

Definition 3.1 (Thurston metric)**.** Define $d_{Lip} = d_{Thurston}$ as (the log of) the optimal Lipschitz[8] constant (in the right homotopy class), i.e.

$$d_{Thurston}(S_1, S_2) = \inf \{\log k : \exists f : S_1 \to S_2 \text{ } k\text{-Lipschitz}, f \circ h_1 \simeq h_2\}.$$

We remark that this was initially called the Lipschitz metric, but at some point people decided that Thurston's name should be attached to it, since his seminal work established so much of the theory and continues to inspire its development.

Proposition 3.2. *This is a weak (i.e. asymmetric) metric.*

Proof. This is actually somewhat easier than the corresponding result for the Teichmüller metric:

- $d_{Lip}(S_1, S_3) \leq d_{Lip}(S_1, S_2) + d_{Lip}(S_2, S_3)$ follows from $\mathrm{Lip}(f_2 \circ f_1) \leq \mathrm{Lip}(f_2)\,\mathrm{Lip}(f_1)$
- $d_{Lip}(S_1, S_2) \geq 0$ with equality iff $S_1 = S_2$ follows from Gauß-Bonnet and length rigidity.
- $d_{Lip}(S_1, S_2) = d_{Lip}(S_2, S_1)$ is not true in general: e.g. consider two genus-2 hyperbolic surfaces, one with a short curve (of say length a) and the other with a even shorter curve (of say length $b \ll a$. By the collar lemma, these short curves have cylinder neighborhoods of width $\sim |\log a|$ and $\sim |\log b| \gg\sim |\log a|$ resp. Suppose these two surfaces have very similar metrics outside of these cylinder neighborhoods.

  Then the optimal Lipschitz constant going in one direction is approximately $\frac{a}{b}$, but going in the other it will be about $\frac{\log b}{\log a}$, which could be much smaller ...

□

We could symmetrize this to get a *bona fide* metric, but the weak metric has a better geomeric interpretation:

Definition 3.3. Given two hyperbolic metrics S_1, S_2 on Σ_g, define

$$K(S_1, S_2) = \sup_{\alpha} \log \frac{\ell_{S_1}(\alpha)}{\ell_{S_2}(\alpha)}$$

where the sup is taken over all closed curves α on Σ_g.

Theorem 3.4 (Thurston)**.** $K(S_1, S_2) = d_{Lip}(S_1, S_2)$

Idea of proof. The $\leq$ direction is obvious (the Lipschitz constant is bounded below by the stretching we see along closed curves.)

For the $\geq$ direction, we explicitly build Lipschitz homeomorphisms ("stretch maps") $f : S_1 \to S_2$ with Lipschitz constant $K(S_1, S_2)$.

This starts with maps between ideal triangles with well-controlled Lipschitz constants

and proceeds by gluing these up along an ideal triangulation ("maximally-stretched laminations"—also see below.)

Within the triangles, the Lipschitz constant of the result is controlled by ideal triangle maps we start with; however it takes considerable technical work to check that one does not pick up more stretch going across the triangulation; this forms a substantial chunk of Thurston's paper. □

[8] Given two (homeomorphic) hyperbolic surfaces S_1, S_2, a homeomorphism $f : S_1 \to S_2$ is k-Lipschitz if for all $p, q \in S_1$ with $p \neq q$, we have

$$\frac{d_{S_2}(f(p), f(q))}{d_{S_1}(p, q)} \leq k.$$

Given $f : S_1 \to S_2$, let $\mathrm{Lip}(f)$ denote the inf of all k such that the above inequality holds for all $p, q \in S_1$ with $p \neq q$.

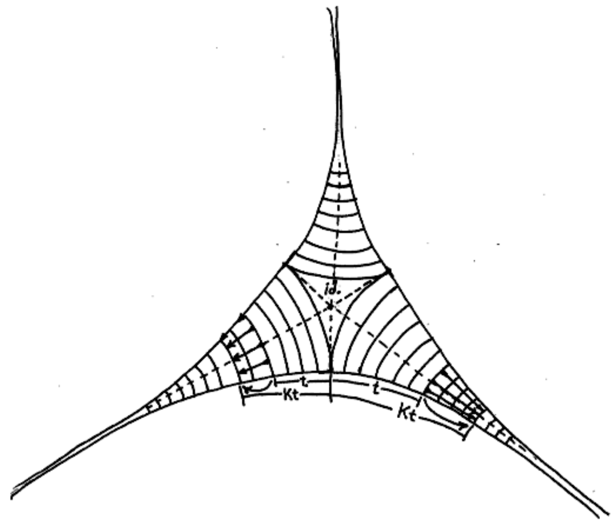

Figure 4: “For each $K > 1$ there is a K-Lipschitz map of an ideal triangle to itself which exactly expands each side by K” (figure from Thurston.)

Some properties of d_{Teich}:

- Finsler, not Riemannian (but not inner products)
- complete
- geodesic but not uniquely geodesic

  Thurston described some of the geodesics (“stretch lines”), but these are not all of them.
- coarsely equivalent to the Teichmüller metric

  (in particular, not Gromov-hyperbolic / negatively-curved in any “obvious” way, but exhibits more subtle features of negative / non-positive curvature.)
- Mod_g-invariant

We remark that Thurston’s results have been generalized to the setting of geometrically-finite hyperbolic manifolds in all dimensions by Guéritaud-Kassel (2016), and may be generalized to other settings ...

3.2 Shearing coordinates

The large chunk of work that Thurston did in studying stretch maps and maximally stretched laminations (see idea of proof for Theorem 3.4) also leads to another set of coordinates on Teichmüller space, which Thurston called shear / cataclysm coordinates.

(The two are supposed to be slightly different ways of conveying the same information, but the difference seems to often get blurred / the names confused in the literature ...)

The idea behind these coordinates is this: start with an ideal triangulation of our closed surface. A pair of pants can divided into two ideal triangles, using a “spinning” triangulation:

[TO INSERT: illustration of spinning triangulation]

Given a pair of hyperbolic ideal triangles which share a single edge in common, we define the shear between them as follows: drop perpendiculars from the two remaining vertices onto the common edge; the (signed) distance between two footpoints is the shear. (For an illustration, see Figure 8 below.)

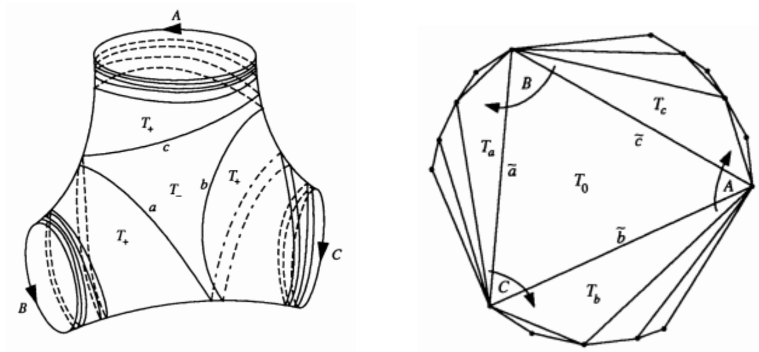

Figure 5: A triangulation of a pair of pants using two ideal triangles—as seen on the pants (left) and in [part of] the universal cover (right) (figure from Goldman.)

To prove that these actually give coordinates, Thurston used **train tracks**, i.e. neighborhoods of geodesic laminations (ideal triangulations) together with measured foliations of these neighborhoods (these are "thick train tracks"; there are also "thin train tracks", which are graphs decorated with numbers which capture the same information. The two notions are meant to be equivalent; the latter highlights the combinatorial aspects of train tracks, the former their geometric / topological origins.)

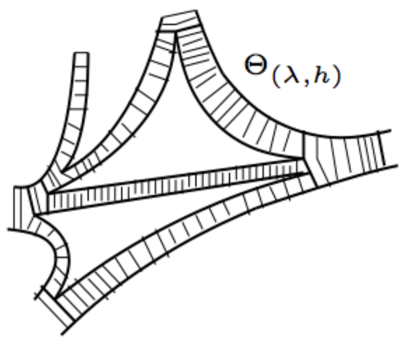

Figure 6: Part of a thick train track (figure from Théret.) Thin train tracks are described further in the Wikipedia article on train tracks [as "train tracks with weights".]

For more details, we refer the reader to Thurston's "Minimal stretch maps between hyperbolic surfaces".

4 Metrics from complex geometry

(Note: this was not talked about during the minicourse.)

We remark, very briefly, that Teichmüller space can be given the structure of a complex manifold in a natural way, e.g. via the Bers embedding into the space of holomorphic quadratic differentials (see e.g. §6.1 of Imayoshi-Taniguchi.)

In fact, the Bers embedding realizes Teichmüller space as a (Kobayashi-hyperbolic) bounded domain of holomorphy.

This allows us to use various general constructions for metrics coming from complex geometry: the Bergman metric, Carathéodory metric, Kähler-Einstein metric, and the Kobayashi metric.

5 Penner's lambda-lengths

Now we consider a different coordinate system, on the Teichmüller space of a *punctured* (hyperbolizable) surface.

Let Σ_g^s be a hyperbolizable (i.e $\chi(\Sigma_g^s) = 2-2g-s < 0$) surface of genus g with $s > 0$ punctures (the theory also works for surfaces with r boundary components and possibly $s = 0$ punctures, as long as $r + s > 0$, but we omit this here for simplicity.)

As above, we define $\mathcal{T}(\Sigma_g^s)$ to be the space of hyperbolic metrics on Σ_g^s, up to isotopy, or equivalently the space of conformal structures on Σ_g^s up to isotopy.

Given Σ_g^s, we can fix an ideal triangulation μ (now without any of this "spinning" trickery, the edges simply go between punctures.)

Given a hyperbolic metric on Σ_g^s, we "decorate" (i.e. associate / assign to) each puncture a horocycle. Then, for each arc α in the ideal triangulation, we define $\delta(\alpha)$ to be the hyperbolic length of α truncated by the horocycles at both ends, and the λ-length of α to be $e^{\delta(\alpha)/2}$.

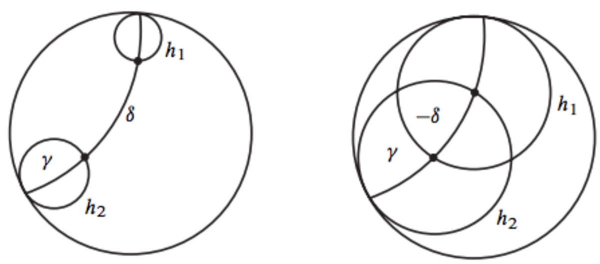

Figure 7: The λ-length associated to the pair of horocycles h_1 and h_2 is $e^{\delta/2}$, where $\delta > 0$ is the [signed] distance indicated above (figure from Penner.)

Definition 5.1. The **decorated Teichmüller space** $\widetilde{\mathcal{T}}(\Sigma_g^s)$ is a trivial $\mathbb{R}^s_{>0}$-bundle over $\mathcal{T}(\Sigma_g)$ (where the fibre should be interpreted as choices of decorations, i.e. horocycles, at each puncture.)

(It can be given a topology by giving $\mathcal{T}(\Sigma_g^s)$ a topology in any of the ways discussed above, and then taking the product topology on $\mathcal{T}(\Sigma_g^s) \times \mathbb{R}$.)

Theorem 5.2 (Penner)**.** *Let Δ be (the set of geodesic arcs in) an ideal triangulation of Σ_g^s.*

Then the map $\widetilde{\mathcal{T}}(\Sigma_g^s) \to \mathbb{R}^{\Delta}_{>0}$ given by sending a hyperbolic metric to the associated set of λ-lengths on Δ is a homeomorphism.

Remark 5.3. Fenchel-Nielsen coordinates play well with the Weil-Petersson symplectic form, but don't transform well under the Mod_g-action.

Fricke coordinates (see below) transform well under the Mod_g-action, but don't play well with the Weil-Petersson symplectic form.

λ-lengths play well with both.

Remark 5.4. λ-lengths are related to Thurston's shear coordinates:

If a, b, c, d are the λ lengths of a ideal quadrilateral, made out of triangle (I) with the sides of length a, b and one of the diagonals, and triangle (II) with sides of length c, d and the same diagonal, then the shear between the two triangles is given by $\log \frac{bd}{ac}$ (for any choice of decorations on the four punctures forming the vertices of the quadrilateral.)

Remark 5.5. λ-lengths depend on a choice of ideal triangulation μ, but the dependence on μ has good properties: the transformations betwene λ-lengths corresponding to different μ faithfully describe the action of the mapping class group on Σ_g^s, and are explicitly calculable in terms of Ptolemy transformations, leading to connections to cluster algebras, quantum Teichmüller theory, and so on ...

Remark 5.6. These coordinates tropicalize to nice coordinates on Thurston's boundary for Teichm uller space.

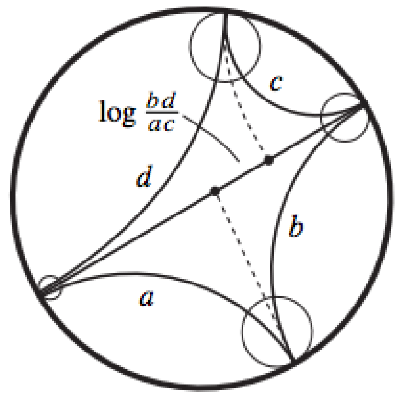

Figure 8: A shear coordinate, with relation to λ-lengths (figure from Penner.)

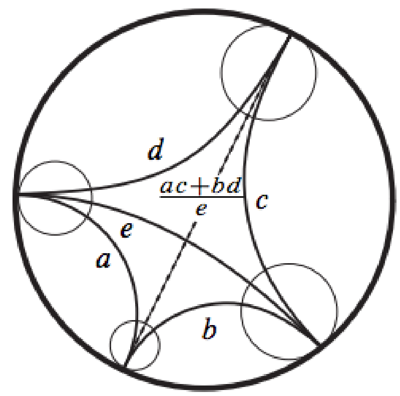

Figure 9: A Ptolemy transformation (figure from Penner.)

Remark 5.7. The theory of λ-lengths can be greatly generalized (see e.g. Fock-Goncharov's work), and has applications to 3-manifolds, string theory, and so on ...

For further details, we refer the reader to Robert Penner's excellent monograph, "Decorated Teichmüller Theory". (A fun excerpt from the Epilogue: "Though it is pooh-poohed by the poo-bahs, we remain boyishly optimistic that the punctured solenoid may be related to Grothendieck's absolute Galois theory beyond the obvious adelic coincidence, perhaps involving some pronilpotent analogue.")

Before finishing with some applications, we digress briefly to talk about the algebraic topology and Fricke coordinates.

A hyperbolic metric on a surface Σ_g^s can be described in terms of the (holonomy) representation it induces of the fundamental group $\pi_1(\Sigma_g^s)$ into the isometry group $\mathrm{PSL}(2,\mathbb{R})$ of $\mathbb{H}^2$.

From this point of view, Teichmüller space can be thought of as a space of representations (sitting inside) $\mathrm{Hom}(\pi_1(\Sigma_g^s), \mathrm{PSL}(2,\mathbb{R}))/\,\mathrm{PSL}(2,\mathbb{R})$ (where the quotient is by the action of $\mathrm{PSL}(2,\mathbb{R})$ by conjugation, which corresponds to the quotient by isotopy in the definition of Teichmüller space.)

$\mathrm{Hom}(\pi_1(\Sigma_g^s), \mathrm{PSL}(2,\mathbb{R}))/\,\mathrm{PSL}(2,\mathbb{R})$ has a natural topology, induced from the topology of $\mathrm{PSL}(2,\mathbb{R})$ as a Lie group (since $\pi_1(\Sigma_g^s)$ is a finitely-presented group and the quotient is nice enough, at least locally ...) This is what is called **the algebraic topology** (unrelated to the field of algebraic topology) on Teichmüller space.

Viewing Teichmüller space as this representation variety also gives us another way of putting coordinates on it:

Definition 5.8 (Fricke coordinates)**.** A homeo $\mathrm{Hom}(\pi_1(\Sigma_g^s), \mathrm{PSL}(2,\mathbb{R}))/\mathrm{PSL}(2,\mathbb{R}) \to \mathbb{R}^{6g-6+s}$ given by

$$[\rho] \mapsto (a_1, b_1, c_1, \ldots, a_{g-1}, b_{g-1}, c_{g-1}, a_1', b_1', c_1', \ldots, a_{g-1}', b_{g-1}', c_{g-1}', t_1, \ldots, t_s)$$

where

- $\pi_1(\Sigma_g) := \langle \alpha_1, \ldots, \alpha_g, \beta_1, \ldots \beta_g, \gamma_1, \ldots, \gamma_s | [\alpha_1, \beta_1] \cdots [\alpha_g, \beta_g] = 1 \rangle$,
- $\rho(\alpha_i) = \begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix}$, $\rho(\beta_i) = \begin{pmatrix} a_i' & b_i' \\ c_i' & d_i' \end{pmatrix}$, and we choose the conjugation by $\mathrm{PSL}(2,\mathbb{R})$ so that β_g is diagonal with ∞ as the attracting fixed point (in the upper-half plane model) and α_g has attracting fixed point at 1.
- $\rho(\gamma_i)$ is conjugate to $\begin{pmatrix} 1 & t_i \\ 0 & 1 \end{pmatrix}$ (the choice of conjugating element will have been determined by the normalization in the previous step.)

A quick dimension count checks out: there are $2g$ non-parabolic generators, each of which contributes 3 degrees of freedom (since their images are in PSL); there is one relation which cuts down 3 degrees, and the conjugation cuts another 3 degrees, for a total of $6g-6$. There are s parabolic generators, each of which contributes 1 degree of freedom, for a total now of $6g-6+s$.

6 A potpourri of applications

6.1 The Nielsen-Thurston classification of mapping classes

(This will be an application of all / any of the three metrics described above.)

Elements in $\mathrm{PSL}(2,\mathbb{R})$ are exactly one of three types:

- elliptic (fix a point $p \in \mathbb{H}^2$ and rotate around it)
- parabolic (fix a point $\xi \in \partial\mathbb{H}^2$, and preserve horocycles based at ξ.)
- hyperbolic (fix two points $\xi, \eta \in \partial\mathbb{H}^2$ and translate along the geodesic axis between them)

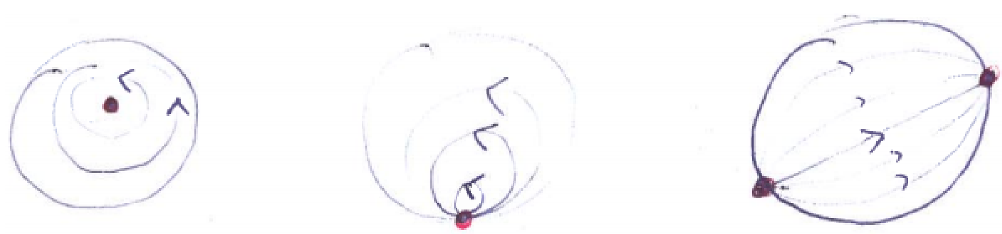

Figure 10: Schematic diagrams of action of elliptic, parabolic, and hyperbolic elements on Poincaré disk; large reddish dots indicate fixed points.

We can characterize these in terms of the **translation distance** τ

Definition 6.1. For $\gamma \in \mathrm{PSL}(2,\mathbb{R})$, $\tau(\gamma) := \inf_{x \in \mathbb{H}^2} d(x, \gamma \cdot x)$

Elliptics have $\tau = 0$ and achieve this inf; parabolics have $\tau = 0$ but do not achieve this inf; hyperbolics have $\tau > 0$ and achieve this inf. (Note that, in the case of $\mathrm{PSL}(2,\mathbb{R})$ and hyperbolic space, we do not have isometries which have positive τ but do not attain it.)

We can use this notion of translation distance to characterize elements of the mapping class group Mod_g. To do so, we will need to replace $\mathbb{H}^2$ by some space on which Mod_g naturally acts by isometries. This space will be $\mathcal{T}(\Sigma_g)$, with any of the metrics (Teichmüller, Weil-Petersson, Thurston) discussed above:

Definition 6.2. For $\phi \in \mathrm{Mod}_g$, $\tau(\phi) := \inf_{S \in \mathcal{T}(\Sigma_g)} d(S, \phi \cdot S)$

Then we define a mapping class group to be ...

- elliptic if it has $\tau = 0$ and achieves this
- parabolic if it does not achieve its translation distance τ (in this case we can have zero *or* positive τ.)
- hyperbolic if it has $\tau > 0$ and achieves this

The Nielsen-Thurston classification (as proved by Bers, using τ) is then (the sum total of) the following series of results:

- Elliptic mapping class groups are periodic (i.e. of finite order.)
- Hyperbolic mapping class groups are pseudo-Anosov (i.e. they preserve a pair of transverse singular foliations, and stretch / contract along them.)
- Parabolic elements are reducible (i.e. they fix some closed curve.)
- Pseudo-Anosov mapping classes are neither reducible nor periodic.

Applying this argument inductively, we find that a general mapping class ϕ will fix some system of closed curves; if we choose this system to be maximal, then restricted to each of the complementary regions ϕ acts as (is isotopic to) either a pseudo-Anosov or a periodic mapping class.

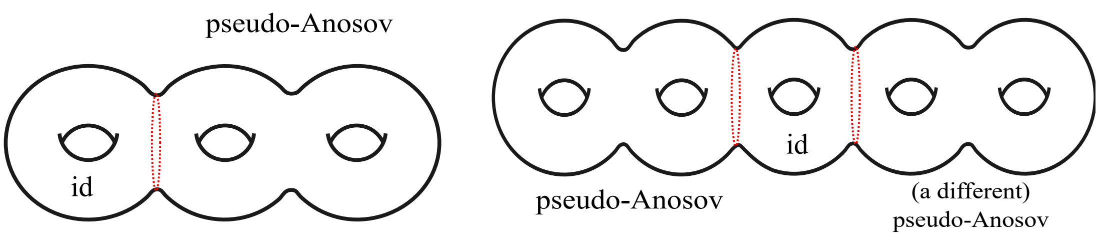

Figure 11: Some examples of general mapping classes—these are divided by multicurves (systems of closed curves) into pieces which (after passing to a power) are either the identity, or pseudo-Anosov.

Example 6.3. A Dehn twist about a separating closed curve C fixes C, and is isotopic to the identity on each of the complementary components.

A Dehn twist about a non-separating closed curve C fixes C, and is isotopic to the identity on the (unique) complementary component.

For details we refer the reader to e.g. §13 of Farb-Margalit.

6.2 Projective embedding of moduli space

We remark that the Weil-Petersson metric (and its properties) can be used to construct an embedding of $\mathcal{M}_g$ into (complex) projective space. For details, we refer the reader to §5.3 of John Harer's notes on "The Cohomology of Moduli Spaces of Curves" (note that pp. 188 and 189 are swapped in those notes, i.e. one should go from page 187 to 189 to 188 to 190.)

6.3 Counting curve types

(This is an application of the Thurston metric.)

Definition 6.4. Define $\mathcal{N}_{k,g}$ to be the number of Mod_g-orbits of closed curves $\gamma \subset \Gamma_g$ with self-intersection $i(\gamma, \gamma) \leq k$.

A Mod_g-orbit of (homotopically non-trivial) closed curves is also called a "curve type". e.g. on the genus-2 surface there are exactly two curve types with zero self-intersection: non-separating simple curves and separating simple curves.

Theorem 6.5 (Aougab-Souto 2017)**.** *Given any $\delta > 0$, $\exists k = k(delta, g)$ s.t. for all $T \geq k$,*

$$e^{\pi\sqrt{|\chi(S)|T}-\delta} \leq \mathcal{N}_{T,g} \leq e^{4\sqrt{2|\chi(S)|T}+\delta}$$

Idea of proof. For the **lower bound**: we combine the following ingredients. Fix a hyperbolic metric S on Σ_g. Let $\mathcal{C}_L(S)$ denote the set of closed geodesics on S of length at most L.

- (Lalley) $\forall\delta > 0$, we have

$$\lim_{L\to\infty} \frac{|\left\{\gamma \in \mathcal{C}_L(S) : \left|\frac{i(\gamma,\gamma)}{\ell_S(\gamma)^2} - \frac{1}{\pi\chi(S)}\right| \leq \delta\right\}}{|\mathcal{C}_L(S)|} = 1$$

  (Roughly, $i(\gamma, \gamma)$, for a random γ, is essentially $\propto \ell_S(\gamma)^2$.)

- (Huber) $|\mathcal{C}_L(S)| \sim \frac{e^L}{L}$

  (Together with Lalley's result, this says that $|\mathcal{C}_{\sqrt{T}}(S) \sim e^{\sqrt{T}}$ and almost all of the curves γ in this set have $i(\gamma, \gamma) \leq T$.)

- Polynomial upper bounds on $|\,\mathrm{Mod}_g \cdot\gamma \cap \mathcal{C}_{\sqrt{T}}(S)|$ obtained by looking at $\mathcal{T}(\Sigma_g)$ with the Thurston metric.

For the **upper bound**:

1. $i(\gamma, \gamma) \leq T \implies \exists$ a hyperbolic metric S on Σ_g s.t. $\ell_S(\gamma) \leq 4\sqrt{2|\chi(S)|T}$ and $\mathrm{injrad}(S) > e^{-4\sqrt{2|\chi(S)|T}}$ where injrad denotes the injectivity radius (so it suffices to look at some ϵ-thick part of moduli space.)

   (The proof of this is kind of cool and slightly magical, but as it is rather off-tangent here we will not explain further. See the Aougab-Suoto paper.)

2. (Given $\epsilon > 0$, we can find δ s.t.) $\exists$ a δ-net of the ϵ-thick part of moduli space with the (symmetrized) Thurston metric of size $\leq f(\delta)|\log(\epsilon)|^N$, where N and f depend only on S.

3. If two points $X, Y \in \mathcal{T}(\Sigma_g)$ are $\leq c$ apart in the symmetrized Thurston metric, then for any closed curve γ, $\ell_Y(\gamma) \leq e^c \ell_X(\gamma)$.

   This makes use of Theorem 3.4, which is true in general only for the unsymmetrized Thurston metric, but some (slightly coarsified) version of it is true when c is taken to be sufficiently small, and this version suffices to give us the bound we want here.

Hence any curve which contributes to $\mathcal{N}_{T,g}$ shows up as a short curve at some point in the thick part of moduli space, and we can approximate [short curves on] this thick part of moduli space by just looking at points in our δ-net. Then we bound $\mathcal{N}_{T,g}$ from above by the total number of points in our δ-net, multiplied by the total number of short curves on each of these points in moduli space (using the results described earlier.) □

For full details we refer the interested reader to Aougab and Souto's paper.

6.4 Computational biology

The combinatorial framework and thinking underlying the theory of λ-lengths leads to useful combinatorial models for certain central problems in computational biology.

One of these problems is **RNA folding**: given the primary structure of a RNA molecule, i.e. the sequence of nucleotides in the RNA (which we may think of as a word in 4 letters), can we compute the free-energy-minimizing configuration(s) of the molecule?

One can build a chord diagram, starting with the primary structure laid out along a line, and with chords for hydrogen bonds that form between nucleotides in the folded configuration (and are the primary mechanism for effecting the folding.)

One can then inflate the chord diagram into a fatgraph (i.e. take some sort of metric neighborhood), and now given a finite fatgraph there is canonical way of associating a punctured Riemann surface, so that it is possible to define a "moduli space of RNA shapes" R_g which is related to the more familiar (?) moduli space of (punctured) Riemann surfaces.

Something similar works for the thornier problem of **protein folding**: here the primary structure is the sequence of amino acids along the protein backbone (a word in twenty letters), there is some folding that is effected hydrogen bonds (the secondary structure), but there is also an additional level of complication (the tertiary structure) due to the fact that amino acids are themselves fairly large molecules, and their constituent atoms have additional degrees of freedom.

The construction of fatgraphs and surfaces starting with chord diagrams in this case consequently involves further constraints, coming from the geometry of these constituent atoms and their interaction/s.

For more details, see §6.4 of Penner's book and the references therein.

References

(Roughly in order of appearance)

- Benson Farb and Dan Margalit, *A Primer on Mapping Class Groups.* Princeton University Press (2012).
- Yoichi Imayoshi and Masahiko Taniguchi, *An Introduction to Teichmüller Spaces.* Springer-Verlag (1992).
- John Harer, "The Cohomology of Moduli Spaces of Curves". In *Theory of moduli (Montecatini Terme, 1985)*, pp. 138–221. Springer (1988), LNMS 1337.
- Francis Bonahon, "The geometry of Teichmüller space via geodesic currents".
- Martin Bridgeman, Richard Canary, and Andrés Sambarino, "An introduction to pressure metrics on higher Teichmüller spaces". arXiv:1507.03894
- Curtis McMullen, "Thermodynamics, dimension and the Weil-Petersson metric".
- Robert Penner, *Decorated Teichmüller Space.* EMS (2012), QGM Master Class series.
- Tarik Aougab and Juan Souto, "Counting curve types". arXiv:1606.06067