

Mapping Class Groups
MSRI, Fall 2007
Days 4, 5, 6,
September 27, October 4 and 11

Lectures and slides by Lee Mosher
Additional notes and diagrams by Yael Algom Kfir

November 20, 2007

Last time:

Theorem 1 (Uniqueness of pseudo-Anosov homeomorphisms). *Each pseudo-Anosov mapping class on S is represented by pseudo-Anosov homeomorphisms uniquely up to topological conjugacy.*

To be precise: given $\Phi_0, \Phi_1 \in \text{Homeo}_+(S)$ pseudo-Anosov homeomorphisms in the same mapping class, there exists $\Psi \in \text{Homeo}_0(S)$ such that

$$\Phi_1 = \Psi\Phi_0\Psi^{-1}$$

Note that the uniqueness is up to “conjugacy by isotopy”, meaning that the topological conjugating map Ψ is isotopic to the identity.

The proof given here is valid only for the case when S is closed. The general case, when S has punctures, requires dealing with the Nielsen theoretic behavior of cusps.

Step 1: Leaves of $\tilde{\mathcal{F}}^s, \tilde{\mathcal{F}}^u$ in $\tilde{S} \approx \mathbf{H}^2$ are uniformly quasigeodesic, because they are geodesic in

$$\mathbf{E}_s^2 = \tilde{S} \text{ with singular Euclidean metric}$$

which is CAT(0) and quasi-isometric to \mathbf{H}^2 .

Notation for Steps 2 and 3 (changed from last time):

- Fix any hyperbolic structure on S ,
- $\Phi =$ any pseudo-Anosov homeomorphism of S ,
- $\tilde{\Phi}: \mathbf{E}_s^2 \rightarrow \mathbf{E}_s^2$ a lift of Φ
- $\Psi =$ any homeomorphism of S isotopic to Φ ,

- $\tilde{\Psi} : \mathbf{H}^2 \rightarrow \mathbf{H}^2$ the lift of Ψ such that

$$\tilde{\Phi} \mid \partial\tilde{S} = \tilde{\Psi} \mid \partial\tilde{S}$$

An argument for the existence of such lifts of Φ and Ψ : Suppose $h : I \times S \rightarrow S$ is an isotopy from Φ to Ψ . Let $p : \mathbf{E}_s^2 \rightarrow S$ be the universal covering space for S . By covering space theory, we can lift $\Phi \circ p$ to $\tilde{\Phi} : \mathbf{E}_s^2 \rightarrow \mathbf{E}_s^2$. Moreover, h lifts to a homotopy $\tilde{h} : \mathbf{E}_s^2 \rightarrow \mathbf{E}_s^2$ so that $\tilde{h}(0, x) = \tilde{\Phi}(x)$. This lifted homotopy is bounded, by the following argument. $\pi_1(S)$ acts by covering transformations cocompactly on \mathbf{E}_s^2 and a fundamental domain F for the action has diameter $\leq 2\text{Diam}(S)$. For any $g \in \pi_1(S)$, since $p \circ \tilde{h}(t, g \cdot x) = h(t, p(x))$, there exists a g' s.t. $\tilde{h}(t, g \cdot x) = g' \cdot \tilde{h}(t, x)$. Now for any $x \in \mathbf{E}_s^2$ find g so that $g^{-1}x \in F$, hence

$$\begin{aligned} d(\tilde{h}(t, x), \tilde{h}(s, x)) &= d(\tilde{h}(t, gg^{-1}x), \tilde{h}(s, gg^{-1}x)) \\ &= d(g'\tilde{h}(t, g^{-1}x), g'\tilde{h}(s, g^{-1}x)) \\ &= d(\tilde{h}(t, g^{-1}x), \tilde{h}(s, g^{-1}x)) \\ &\leq 2D \end{aligned}$$

In particular $\tilde{\Phi}$ and $\tilde{\Psi}$ induce the same action on the boundary. \diamond

This proof shows that not only do $\tilde{\Phi}, \tilde{\Psi}$ have the same action on the circle at infinity, but more strongly their values in the finite plane differ by at most a constant; in what follows we shall denote the bounding constant by C_1 .

Step 2: $\tilde{\Psi}$ contracts \mathbf{H}^2 geodesics that have the same endpoints in $\partial\tilde{S}$ as leaves of $\tilde{\mathcal{F}}^s$. To be precise:

There exists $A \geq 0$ such that if

- $\ell_0^e = \text{leaf of } \tilde{\mathcal{F}}^s$
- $\ell_0 = \mathbf{H}^2$ geodesic with same endpoints as ℓ_0^e in $\partial\tilde{S}$
- $\ell_i^e = \tilde{\Phi}(\ell_{i-1}^e)$ (inductive definition)
- $\ell_i = \text{straightening of } \tilde{\Psi}(\ell_{i-1})$ (inductive definition)
= \mathbf{H}^2 geodesic with same endpoints as ℓ_i^e in $\partial\tilde{S}$

then for all $x_0, y_0 \in \ell_0$, letting

- $x_i, y_i \in \ell_i$ be closest to $\tilde{\Psi}(x_{i-1}), \tilde{\Psi}(y_{i-1}) \in \tilde{\Psi}(\ell_{i-1})$
(inductive definition)

it follows that there exists N such that if $i \geq N$ then

$$d_{\mathbf{H}^2}(x_i, y_i) \leq A$$




Figure 1: The straight (green) lines are geodesics in \mathbf{E}_s^2 the curvy (red) lines are geodesics in \mathbf{H}^2 . Since these metrics are q.i. to each other, and $\tilde{\Phi}$ and $\tilde{\Psi}$ are a bounded distance apart, ℓ_i^e and ℓ_i define the same endpoints on the boundary with respect to either metric.

Some constants for the proof of Step 2

λ = expansion factor for Φ
 = bilipschitz constant for $\tilde{\Phi}: \mathbf{E}_s^2 \rightarrow \mathbf{E}_s^2$
 K, C = quasi-isometry constants for maps

$$\tilde{\Psi}: \mathbf{H}^2 \rightarrow \mathbf{H}^2 \quad \text{and} \quad \text{Id}: \mathbf{H}^2 \leftrightarrow \mathbf{E}_s^2$$

C_1 = closeness constant for straightening
 K, C quasigeodesics to geodesics
 in \mathbf{H}^2 and in \mathbf{E}_s^2
 = closeness constant for maps $\tilde{\Phi}, \tilde{\Psi}$ on \mathbf{E}_s^2

For example, $d(\tilde{\Psi}(x_{i-1}), x_i) \leq C_1$.

Strategy of the proof of Step 2: Iterate in \mathbf{E}_s^2

Switching from \mathbf{H}^2 to \mathbf{E}_s^2 . Letting $x_i^e =$ point on ℓ_i^e closest to x_i in \mathbf{E}_s^2 ,

$$d_{\mathbf{E}_s^2}(x_i, x_i^e) \leq C_1$$

which implies

$$d_{\mathbf{H}^2}(x_i, x_i^e) \leq KC_1 + C_0$$

Similarly for y_i^e .

Step 2 therefore reduces to proving:

- The sequence x_i^e, y_i^e is contracted in \mathbf{E}_s^2 , i.e.
- $d_{\mathbf{E}_s^2}(x_i^e, y_i^e)$ eventually falls below some constant.

Iteration in \mathbf{E}_s^2

$$d(x_i^e, \tilde{\Phi}(x_{i-1}^e)) \leq (\lambda + K + 2)C_1 + C_0 = C_2$$

map by $\tilde{\Phi}$ and make C_2 error;
map by $\tilde{\Phi}$ and make C_2 error;
map by $\tilde{\Phi}$ and make C_2 error;
 \vdots

$$\begin{aligned} d(x_i^e, y_i^e) &\leq d(x_i^e, \tilde{\Phi}(x_{i-1}^e)) \\ &\quad + d(\tilde{\Phi}(x_{i-1}^e), \tilde{\Phi}(y_{i-1}^e)) \\ &\quad + d(\tilde{\Phi}(x_{i-1}^e), y_i^e) \\ &\leq \frac{1}{\lambda}d(x_{i-1}^e, y_{i-1}^e) + 2C_2 \end{aligned}$$

shrink by λ and make $2C_2$ error;
shrink by λ and make $2C_2$ error;
shrink by λ and make $2C_2$ error;
 \vdots

It follows that $d(x_i^e, y_i^e)$ eventually shrinks below $\frac{2C_2}{1 - \frac{1}{\lambda}}$.

Explicitly: After the n -th iteration

$$\begin{aligned} d(x_n^e, y_n^e) &\leq \frac{1}{\lambda^n}d(x_0, y_0) + C_2\left(1 + \frac{1}{\lambda} + \frac{1}{\lambda^2} + \cdots + \frac{1}{\lambda^{n-1}}\right) \\ &\leq 1 + \frac{2C_2}{1 - \frac{1}{\lambda}} \end{aligned}$$




Figure 2: x_1^e and x_1 are C_1 close in the Euclidean metric. x_1 and $\tilde{\Psi}(x_0)$ are C_1 close in the hyperbolic metric so $KC_1 + C$ close in the Euclidean metric. $\tilde{\Phi}(x_0)$ and $\tilde{\Psi}(x_0)$ are C_1 close in the hyperbolic metric, so $KC_1 + C$ close in the Euclidean metric. $\tilde{\Phi}(x_0)$ and $\tilde{\Phi}(x_0^e)$ are λC_1 close because x_0 and x_0^e are C_1 close and Φ stretches distances by at most λ . In conclusion, x_1^e and $\tilde{\Phi}(x_0^e)$ are C_2 close.

What makes this calculation work is that the constant C_2 does not depend on the step i that we're at.

◇ Step 2

Step 3: $\tilde{\Psi}$ does *NOT* contract any \mathbf{H}^2 geodesic that does *NOT* have the same endpoints in $\partial\tilde{S}$ as any leaf of $\tilde{\mathcal{F}}^s$.

To be precise, given:

- $\ell_0 = \mathbf{H}^2$ geodesic
- $\ell^e = \mathbf{E}_s^2$ geodesic with same endpoints
- ℓ_i, ℓ_i^e as before (inductive)

Assume ℓ^e is *not* a leaf of $\tilde{\mathcal{F}}^s$.

It follows that ℓ^e has a subsegment that is *not* horizontal, which implies that there exists $\exists i$ and a segment $\alpha \subset \ell_i^e$ which is nearly vertical and very long.

Now iterate the process:

- map by $\tilde{\Phi}$ and make bounded error;
- map by $\tilde{\Phi}$ and make bounded error;
- map by $\tilde{\Phi}$ and make bounded error;
- \vdots

This process, applied to x, y in opposite components of $\ell_i^e - \alpha$, does *not* eventually contract distance.

Final Step: Proof of uniqueness of pseudo-Anosov homeomorphisms

Let Φ_i , $i = 0, 1$, be isotopic pseudo-Anosov homeomorphisms, with stable and unstable foliations $\mathcal{F}_i^s, \mathcal{F}_i^u$.

Choose lifts $\tilde{\Phi}_0, \tilde{\Phi}_1$ to be equal at ∞ .

Steps 2 and 3 together imply that $\tilde{\Phi}_0, \tilde{\Phi}_1$ contract the *exact same set of geodesics*, namely those with endpoints

$$\partial^2 \tilde{\mathcal{F}}_0^s = \partial^2 \tilde{\mathcal{F}}_1^s$$

Also, $\tilde{\Phi}_0^{-1}, \tilde{\Phi}_1^{-1}$ contract the *exact same set of geodesics*, namely those with endpoints

$$\partial^2 \tilde{\mathcal{F}}_0^u = \partial^2 \tilde{\mathcal{F}}_1^u$$

Must construct homeomorphism $F: \tilde{S} \rightarrow \tilde{S}$ such that

- ($\pi_1 S$ -equivariance) F respects orbits of the action of $\pi_1 S$, and so descends to S
- $\tilde{\Phi}_0 F = F \tilde{\Phi}_1$

For $i = 0, 1$, each point $x \in \tilde{S}$ is named by two subsets of $\partial\tilde{S}$:

- $\partial_i^s(x)$ = set of endpoints of leaves of $\tilde{\mathcal{F}}_i^s$ passing through x (at least 2 points, possibly more)
- $\partial_i^u(x)$ = set of endpoints of leaves of $\tilde{\mathcal{F}}_i^u$ passing through x (at least 2 points, possibly more)

Lemma (Basis of well-definedness:). *The set of ordered pairs of subsets*

$$\{(\partial_i^s(x), \partial_i^u(x)) \mid x \in \tilde{S}\}$$

is independent of $i = 0, 1$.

Proof. The set of subsets

$$B_i^s = \{\partial_i^s(x) \mid x \in \tilde{S}\}$$

is independent of $i = 0, 1$: follows from steps 2 and 3.

Given $i = 0, 1$, and given a pair of subsets

$$X \in B_i^s, \quad Y \in B_i^u$$

the following are equivalent:

- $\exists x \in S$ such that $(X, Y) = (\partial_i^s(x), \partial_i^u(x))$
- X, Y link each other in S_∞^1 .

This is independent of i . ◇

Now define $F: \tilde{S} \rightarrow \tilde{S}$: for each $x \in \tilde{F}$ define the point

$$x' = Fx$$

by the requirements

$$\begin{aligned} \partial_1^s(x') &= \partial_0^s(x) \\ \partial_1^u(x') &= \partial_0^u(x) \end{aligned}$$

It is easy to see (or have already checked):

- well-defined bijection
- $\pi_1 S$ equivariance
- set theoretic conjugacy

Must also prove:

- continuity




Figure 3: The endpoints of the solid geodesic are in $\partial_i^s(x)$ and the endpoints of the dashed geodesic are in $\partial_i^u(x)$. They are linked.

A few words about continuity

- Topologize finite subsets of $\partial\tilde{S}$ so that A, B are close if they have subsets of cardinality ≥ 2 that are close in the Hausdorff sense.
- The map

$$\tilde{S} \rightarrow \text{finite subsets of } \partial\tilde{S} \times \text{finite subset of } \partial\tilde{S}$$

defined by

$$x \mapsto (\partial_i^s(x), \partial_i^u(x))$$

is a homeomorphism onto its image.

- Although the full space (finite subsets of $\partial\tilde{S} \times$ finite subset of $\partial\tilde{S}$) is not Hausdorff, the image of the map is Hausdorff, because for each x, y , the two sets $\partial_i^s(x), \partial_i^s(y)$ do *not* link each other, and the two sets $\partial_i^u(x), \partial_i^u(y)$ do *not* link each other.
- The image is independent of $i = 0, 1$.

Remark 1. *This model for pseudo-Anosov homeomorphisms was discovered by Culler-Shalen-Levitt. Is there a similar model for fully irreducible outer automorphisms of F_n ?*

Topological conjugacy invariants of pseudo-Anosov homeomorphisms

- Folk knowledge: the set of train track maps is a complete topological invariant (up to combinatorial equivalence of train track maps)
- Proof is almost obvious.

Advantage:

- Easy to say, easy to prove.

Disadvantage:

- Although it is finitistic ... *Usually* one describes a train track map by writing the image of each branch as a sequence of branches.
- Example: consider $M \in \mathrm{SL}(2, \mathbf{Z})$ given by

$$M = \underbrace{RLRL^2RLR^2L \dots RL^2RLRLRLR^2L}_{k \text{ alternations between } R \text{ and } L}$$

where R or R^2 is chosen randomly,
and L or L^2 is chosen randomly.

- Fact: the image of each generator of \mathbf{Z}^2 under this map has length exponential in k .
- Conclusion: writing out images of branches is *vastly* more inefficient than
 - “RL” notation,
 - equivalently, train track splitting information.

Instead we shall use: **Train track expansions**

Given:

- Pseudo-Anosov homeomorphism $\Phi \in \mathrm{Homeo}_+(S)$
- Stable and unstable foliations $\mathcal{F}^s, \mathcal{F}^u$

Construct **train track approximations** of \mathcal{F}^u by:

- Choosing an unstable leaf segment (or family thereof)
- Use it (them) to decompose S into su-rectangles
- Use the rectangle decomposition to construct the train track.

Construct **splittings** (and other relations) between train tracks by:

- perturbing choice of unstable leaf segment (or family thereof)

Construct **train track expansion** by:

- Defining a parameter space for choice of unstable leaf segment (or family thereof)
- See how train track varies, as parameter varies

The “train track expansion” (up to appropriate equivalence) is a full conjugacy invariant.

Comments about leaf segments

- On torus, xy -isometry group acts transitively
 - $\forall r > 0$ we chose one unstable segment ℓ_r .
 - * any two choices of ℓ_r are equivalent under an xy -isometry, and lead to isotopic train tracks
- On S with $\chi(S) < 0$, xy -isometry group is finite on S , only xy -isometry isotopic to identity is identity.
 - Distinct choices of length r unstable segment are inequivalent under x, y -isometry isotopic to identity.
 - Can lead to nonisotopic train tracks
 - Need (semi)canonical choices of unstable leaf segments

Separatrices

- An *infinite separatrix* of a measured foliation is a half-infinite leaf based at a singularity. (PICTURE)
- There are finitely many singularities, and finitely many infinite separatrices of any measured foliation on S . Bounds depend only on topology of S .
- Indexed sets of infinite separatrices:
 - of \mathcal{F}^u , $\{\ell^i \mid i \in I\}$
 - of \mathcal{F}^s , $\{w^j \mid j \in J\}$

Unstable separatrix systems

- Given $i \in I$ and $s \geq 0$, let
 - ℓ_s^i = initial length s subsegment of the unstable separatrix ℓ^i .
- Parameter space for unstable separatrix systems:

$$\mathcal{S} = [0, +\infty)^I - \mathcal{O}$$




Figure 4: For each unstable leaf emanating from the singularity, allocate a subsegment of length l_i




Figure 5: A separatrix system will determine a rectangle decomposition.

- Notation: we let $\underline{s} = (s_i)_{i \in I}$ denote a typical element of \mathcal{S}
- Given $\underline{s} \in \mathcal{S}$, the set

$$\ell_{\underline{s}} = \cup_{i \in I} \ell_{s_i}^i$$

is called a *separatrix system* of \mathcal{F}^u .

Rectangle decompositions

For each $\underline{s} \in \mathcal{S}_\Phi$, construct a *rectangle decomposition* $\mathcal{R}_{\underline{s}}$:

- The union of the horizontal boundaries of the rectangles in $\mathcal{R}_{\underline{s}}$ equals

$$\ell_{\underline{s}} = \cup_{i \in I} \ell_{s_i}^i$$

Let $\partial \ell_{\underline{s}}$ be the set of nonsingular endpoints of those $\ell_{s_i}^i$ for which $s_i \neq 0$.

- The union of the vertical boundaries of the rectangles in $\mathcal{R}_{\underline{s}}$ equals

$$w_{\underline{s}} = \left(\bigcup_{j \in J} w_{\underline{s}}^j \right) \cup \left(\bigcup_{x \in \partial \ell_{\underline{s}}} w_{\underline{s}}^x \right)$$

where

- $w_{\underline{s}}^j$ = longest initial segment of w^j with interior disjoint from $\ell_{\underline{s}}$
- $w_{\underline{s}}^x$ = longest stable leaf segment containing x with interior intersecting $\ell_{\underline{s}}$ solely at x .
- \mathcal{R}_s = closures of components of $S - (\ell_{\underline{s}} \cup w_{\underline{s}})$
- Each $R \in \mathcal{R}_s$ is a rectangle, foliated in the “horizontal” direction by segments of \mathcal{F}^u , and in the “vertical” direction by segments of \mathcal{F}^s .
- The interior of R is embedded and is disjoint from the boundary of R .
- The boundary of R is immersed but may have some identifications.
- Each “horizontal” component of ∂R maps to $\ell_{\underline{s}}$.
- Each “vertical” component of ∂R maps to $w_{\underline{s}}$.

Train tracks

For each $\underline{s} \in \mathcal{S}_\Phi$, construct a *train track* $\tau_{\underline{s}}$:

- $\tau_{\underline{s}}$ has one *main branch* b_R for each $R \in \mathbf{R}_{\underline{s}}$, an interior horizontal segment of R .
- $\tau_{\underline{s}}$ has *main switches* as follows:




Figure 6: Train track switches.

- one main switch σ^j on each $w_{\underline{s}}^j$, $j \in J$
- one main switch σ^x on each $w_{\underline{s}}^x$, for x in the boundary of each positive length ℓ_s^j
- τ_s has one *secondary switch* and two length zero *secondary branches* for each zero length ℓ_s^j

Structure of the parameter space \mathcal{S}_{Φ}

- For $\underline{s} \in \mathcal{S}_{\Phi}$ let $\mathcal{C}_{\underline{s}} = \{s' \in \mathcal{S}_{\Phi} \mid \tau_{\underline{s}}, \tau_{\underline{s}'}$ are isotopic
- Denote $\mathcal{C}_{\underline{s}} = C_{\tau}$ where $\tau = \tau_{\underline{s}}$
- The collection $\{C_{\tau}\}$ is a decomposition of \mathcal{S}_{Φ} , indexed by isotopy classes of those train tracks of the form $\tau_{\underline{s}}$ for some $\underline{s} \in \mathcal{S}_{\Phi}$.

Theorem.

1. Each $C_{\tau} \subset S \subset \mathbf{R}^I$ is the interior of a compact rectilinear polyhedron in \mathbf{R}^I .
2. $\{\overline{C}_{\tau}\}$ is a CW-decomposition of \mathcal{S}_{Φ} , parameterized by isotopy classes invariant train tracks for positive powers of ϕ .
3. A cell C_{τ} is a face of a cell $C_{\tau'}$ if and only if there is a smooth forest collapse

$$\tau' \xrightarrow{F} \tau$$

where F is a subforest of τ . (PICTURES)

- This map must be homotopic to the identity

- *Smooth arcs must map to smooth arcs (not necessarily by a local injection, though)*

4. *The infinite cyclic subgroup $\langle \phi \rangle$ acts by cellular isomorphisms, with $\phi^k(\overline{\mathcal{C}}_\tau) = \overline{\mathcal{C}}_{\phi^k(\tau)}$.*

The proof of the first item in this theorem requires one to identify, for each cell \mathcal{C}_τ , a finite set of equations and inequalities that defines the polyhedral structure on \mathcal{C}_τ .

The generic case occurs when $s_i > 0$ for each i and when each switch of τ has valence ≤ 3 . This case corresponds to a top dimensional cell \mathcal{C}_τ defined solely by inequalities, no equations. As \underline{s} varies, the horizontal length of a rectangle $R \in \mathbf{R}_{\underline{s}}$, which we denote L_R^h , can vary. But if \underline{s} varies too much, L_R^h can shrink to zero. Corresponding to R there is an inequality which says $L_R^h > 0$. In order to show that this is a linear inequality, one needs an affine expression for L_R^h . This expression involves only two of the coordinates s_i, s_j , chosen so that the point $\partial \ell_{s_i}^i$ lies on one vertical boundary of R and the point $\partial \ell_{s_j}^j$ lies on the other vertical boundary of R . The expression for L_R^h has one of two forms, depending on the orientations of $\ell_{s_i}^i$ and $\ell_{s_j}^j$ relative to R , namely:

Ships travelling in opposite directions $L_R^h = s_i + s_j + (\text{constant})$

Ships travelling in the same direction $L_R^h = s_i - s_j + (\text{constant})$

The work involved in proving that \mathcal{C}_τ is a rectilinear polyhedron is to show that $\underline{s} \in \mathcal{C}_\tau$ if and only if \underline{s} satisfies the collection of inequalities $L_R^h > 0$. The “only if” direction is obvious from the discussion above. The “if” direction requires looking at other affine inequalities that can arise from coincidences of finite separatrices, coincidences that are more general versions of the inequalities $L_R^h > 0$ associated to the combinatorics of $\mathbf{R}_{\underline{s}}$ itself, and showing that each of these other affine inequalities is implied by the inequalities $L_R^h > 0$.

The proof of item 3 in the above theorem can be understood for codimension-1 faces of top dimensional cells by letting a defining inequality $L_R^h > 0$ degenerate to an equation $L_R^h = 0$. This is depicted in the two accompanying figures, in the case of “ships travelling in opposite directions”.




Figure 7: These depict two points on a codimension 1 cell defined by $s^i + s^j = A$ for some constant A . Horizontal lines correspond to segments of $\ell_{\underline{s}}$. Vertical lines correspond to segments of the system of $w_{\underline{s}}$.




Figure 8: The top picture depicts a common face of the two cells depicted on the bottom. The bottom left cell corresponds to $s^i + s^j < A$. The bottom right cell corresponds to $s^i + s^j > A$. Their common face corresponds to $s^i + s^j = A$




Figure 9: A part of the train track complex




Figure 10: In general the codimension of a face of a top dimensional cell equals the number of main branches (rectangles) whose horizontal lengths degenerate to zero.

In the above theorem, cells are labelled by isotopy classes of train tracks.

In order to define conjugacy invariants, we must label cells by the combinatorial type of the train track, not just by its isotopy class.

Combinatorial type of a train track

- Two train tracks $\tau, \tau' \subset S$ are *combinatorially equivalent* if $\exists h \in \text{Homeo}_+(S)$ such that $h(\tau) = \tau'$.
- Such a map h is called a *combinatorial equivalence* between τ and τ' .
- Existence of a combinatorial equivalence is easy to check in terms of gluing diagrams for τ, τ' .
- There are only finitely many gluing diagrams, so there are only finitely many combinatorial equivalence classes.

Remark. An example of combinatorial equivalence. On the genus 2 once punctured surface, consider the case where there are five 3-pronged singularities and at the puncture a 1-pronged singularity. Suppose that the only nonzero coordinate of \underline{s} is the one s_i associated to the unstable separatrix ℓ^i at the 1-pronged singularity. Suppose furthermore that we are in the generic case where the stable segment $w_{\underline{s}}^x$ passing through $x = \partial \ell_{s_i}^i$ does not contain any singularity. Then the train track $\tau_{\underline{s}}$ has one punctured monogon, and five trigons whose boundaries consist entirely of secondary branches. There are exactly 105 combinatorial equivalence classes of train tracks that arise in this manner.

Combinatorial type of a smooth subforest collapse

- Given two smooth subforest collapses

$$f: \tau_0 \xrightarrow{F} \tau_1$$

$$f': \tau'_0 \xrightarrow{F'} \tau'_1$$

- They are *combinatorially equivalent* if there exist combinatorial equivalences $h_0: \tau_0 \rightarrow \tau'_0$ and $h_1: \tau_1 \rightarrow \tau'_1$ such that

$$h_1 f \text{ is homotopic to } f' h_0$$

- Again, existence of a combinatorial equivalence is easy to check in terms of gluing diagrams and how the subforests fit into them.

Theorem. *Two pseudo-Anosov homeomorphisms $\Phi, \Phi': S \rightarrow S$ are topologically conjugate if and only if there exists*




Figure 11: A combinatorial equivalence between two smooth subforest collapses.

- a *CW-isomorphism* $h: \mathcal{S}_\Phi \rightarrow \mathcal{S}_{\Phi'}$
- a *combinatorial equivalence* between corresponding pairs of cells, and between corresponding pairs of face inclusions

such that

Compatibility the cell combinatorial equivalences and the face inclusion combinatorial equivalences are compatible up to homotopy

Equivariance Everything is equivariant with respect to Φ and Φ'

Remark: The compatibility condition is useful and necessary:

- It is useful in the proof, providing various homeomorphisms needed to “change the marking” of S in the course of the proof.
- It is necessary in order to “break finite order symmetries” of train tracks.

Proof.

- Given a conjugacy $\Psi: S \rightarrow S$ such that $\Psi\Phi\Psi^{-1} = \Phi'$, Ψ induces the isomorphism h and all of the combinatorial equivalences.
- Conversely, given all of the data as in the theorem, take Ψ to be the mapping class of the combinatorial equivalence between any pair of cells, use compatibility to prove that Ψ is well-defined, and use equivariance to prove that Ψ conjugates Φ to Φ' .

◇

Why is this finistic?

- Because it can be stated in terms of the finite quotient complex $\mathcal{S}_\Phi/\langle\Phi\rangle$, a “simplex bundle over the circle”. This is analogous to the loop of L’s and R’s.
- Existence of a system of “combinatorial equivalences” between $\mathcal{S}_\Phi/\langle\Phi\rangle$ and $\mathcal{S}_{\Phi'}/\langle\Phi'\rangle$ can be checked efficiently and easily by computer.

Why is this reasonable?

- \mathcal{S}_Φ may seem an intractable object.
- Structure of a “simplex bundle over a line”
- Proper homotopy equivalent to a line
- The quotient complex \mathcal{S}_Φ/Φ is compact, and is homotopy equivalent to the circle.
- Given a metric along the edges of the dual 1-complex, define *cross sectional diameter* at $x \in \mathcal{S}_\Phi$ to be $D_x =$ minimum sum of lengths of dual 1-cells in a subcomplex $K \subset \mathcal{S}_\Phi$ such that $x \in K$ and K separates the two ends
- Conjecture/expectation/hope: cross sectional diameter of \mathcal{S}_Φ has an upper bound which is polynomial in the circumference of \mathcal{S}_Φ/Φ .

Why is this computable?

- Apply Bestvina-Handel “Train tracks for surface homeomorphisms”. This provides an invariant train track, if one exists.
- If pseudo-Anosov, an invariant train track does exist
- If there is an invariant train track, there are algorithms to decide if indeed pseudo-Anosov
- Once having an invariant train track for a pseudo-Anosov, can compute all of $\mathcal{S}_\Phi/\langle\Phi\rangle$ by applying “splitting relations” among train tracks
- Given invariant train track τ and a splitting sequence

$$\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_K = \phi(\tau)$$

should be able to algorithmically construct \mathcal{S}_Φ/Φ in time which is polynomial in K (assuming that the conjecture/expectation/hope is true)