

Uniformly polynomial-time classification of surface homeomorphisms

Filippo Baroni

Abstract

We describe an algorithm that, given two essential curves on a surface S , computes their distance in the curve graph of S , up to multiplicative and additive errors. As an application, we present an algorithm to decide the Nielsen-Thurston type (periodic, reducible, or pseudo-Anosov) of a mapping class of S . The novelty of our algorithms lies in the fact that their running time is polynomial in the size of the input *and* in the complexity of S – say, its Euler characteristic. This is in contrast with previously known algorithms, that run in polynomial time in the size of the input *for any fixed surface S* .

Contents

1	Introduction	2
2	Surfaces, curves, and train tracks	3
2.1	Triangulated surfaces	3
2.2	Curve graph	7
2.3	Combinatorics of train tracks	8
3	Coarse distance algorithm	17
3.1	Agol–Hass–Thurston moves	18
3.2	Elementary properties of AHT sequences	22
3.3	Deep nesting algorithm	24
3.4	Distances for carried curves	28
3.5	Distances in the curve graph	31
4	Nielsen-Thurston classification	37
4.1	Periodic homeomorphisms	37
4.2	Subsurface projection	38
4.3	Stable translation length	41
4.4	Algorithmic classification	44
	Constants	45
	References	46

1 Introduction

It was proved by Masur and Minsky [24] that the curve graph $\mathcal{C}(S)$ of a surface S is hyperbolic, and that a mapping class of S is pseudo-Anosov if and only if it acts loxodromically on $\mathcal{C}(S)$. In this article, we aim to exploit this characterisation to provide an efficient algorithm for recognising pseudo-Anosov surface homeomorphisms. Naturally, the crucial building block of this classification algorithm will be a procedure to (coarsely) compute distances in the curve graph.

There is a wide collection of algorithms in the literature addressing the computation of distances in $\mathcal{C}(S)$ (for instance, [8, 31, 33]) and the Nielsen-Thurston classification of surface homeomorphisms (for instance, [4, 7, 15, 21, 32]). These lists include [5], where Bell and Webb describe, for each *fixed* surface, polynomial-time algorithms to solve the two mentioned problems. In contrast, the main contribution of the present article is the focus on providing algorithms that, while following the template of Bell and Webb, run in polynomial time in the size of the input, *uniformly with respect to the complexity of the surface* – by which we mean, for example, its Euler characteristic.

In order to estimate the distance between two curves in $\mathcal{C}(S)$, we rely on a result of Masur and Minsky [26], stating that a splitting sequence of train tracks corresponds to an unparametrised quasi-geodesic in the curve graph. By applying a combinatorial criterion of Masur and Minsky [24] to a carefully constructed splitting sequence, we are able to find the correct quasi-geodesic parametrisation, thus proving the following.

Theorem. *Let a and b be essential curves on S . There is an algorithm to compute an estimate d of the distance $\text{dist}_{\mathcal{C}(S)}(a, b)$ between a and b in $\mathcal{C}(S)$, such that*

$$d - L_+ \leq \text{dist}_{\mathcal{C}(S)}(a, b) \leq L_\times \cdot d + L_+.$$

The constants L_\times and L_+ are polynomial in the complexity of S , and the running time of the algorithm is polynomial in the complexities of S , a , and b .

We will describe in detail what we mean by complexity of curves in Section 2.1; for now, let us say that the surface S is given as input in the form of a triangulation, and the complexity of a curve is logarithmic in the number of times it intersects the edges of the triangulation. Note that the algorithm in the theorem above does not compute the exact distance between two curves, but rather a reasonably accurate estimate thereof. However, this approximation is enough to detect the Nielsen-Thurston type of a surface homeomorphism f . In fact, for an arbitrary curve a on S , the quantity $\text{dist}_{\mathcal{C}(S)}(a, f^k(a))$ grows linearly in k if f is pseudo-Anosov, and remains bounded if f is periodic or reducible. This statement can be made effective, thanks to the work of Gadre and Tsai [14] and Webb [34], and leads to a proof of the following theorem.

Theorem. *Let f be a homeomorphism of S . There is an algorithm to decide if f is periodic, reducible (and non-periodic), or pseudo-Anosov. The running time of the algorithm is polynomial in the complexities of S and f .*

Again, we will explain our notion of complexity of homeomorphisms in Section 2.1; for now, we can take it to mean the word length in the mapping class group of S with respect to a generating set consisting of “reasonably short” Dehn twists.

Let us emphasise, once more, that the novelty of this algorithm resides in the good dependence of the running time on the complexity of the surface. To better illustrate this point, let us consider the braid group B_n on n strands with the standard generating set

$\{\sigma_1, \dots, \sigma_{n-1}\}$. The following result, which is almost a direct generalisation of Calvez’s [11, Theorem 1], is an easy consequence of the classification theorem above.

Corollary. *There is an algorithm to decide the Nielsen-Thurston type of a braid in polynomial time in its word length and in the number n of strands.*

This article is organised as follows. In Section 2 we set up the definitions and notation we will use throughout the article. Most of this section is dedicated to combinatorial properties of train tracks, which will be crucial for our coarse distance algorithm. We also discuss in detail how to represent surfaces, curves, homeomorphisms, and train tracks in a discrete fashion which is suitable to algorithmic manipulations, as well as the associated notions of complexity. Section 3 follows, with a description of the algorithm to estimate distances in the curve graph. The bulk of this section is devoted to the construction of specific splitting sequences of train tracks and their quasi-geodesic subsequences. Finally, in Section 4, we apply the distance algorithm to the Nielsen-Thurston classification problem.

We have made an effort to keep the constants in our algorithms as explicit as possible. For the sake of convenience, we have collected all of them in a table, which can be found at the end of the article.

2 Surfaces, curves, and train tracks

In this section, we introduce definitions and properties of the objects we will be using in our algorithms, including surfaces, curves, homeomorphisms, and train tracks. These objects need to be fed as input to, manipulated by, and returned as output by algorithmic routines. For this reason, alongside the usual, “theoretical” definitions, we will also provide ways to represent these objects combinatorially, and define complexities of these representations.

As anticipated before, in this article we focus on good dependence of the running times on the complexity of the underlying surface. Therefore, all the constants we hide in the notation will be *universal*. For instance, for real-valued functions F and G , we will say that “ F is $\mathcal{O}(G)$ ” if $F \leq \alpha \cdot G + \beta$ for some universal constants $\alpha, \beta \geq 0$. Similarly, we will write “poly(X)” to signify a universal, albeit implicit, polynomial in the variable X ; by “ Y is polynomial in X_1, \dots, X_n ”, we will mean that

$$Y = \mathcal{O}(\text{poly}(X_1) \cdots \text{poly}(X_n)).$$

2.1 Triangulated surfaces

Basic definitions. In the context of this article, by “surface” we always mean a connected, orientable surface S of finite genus $g \geq 0$, with $p \geq 0$ punctures and empty boundary. We will say that S is *closed* if $p = 0$, and *punctured* otherwise.

By “curve on S ” we always mean a simple closed curve embedded in S . A curve $a \subseteq S$ is *essential* if it does not bound a disc or a once-punctured disc in S . Given two curves a and b on S , we denote their (geometric) intersection number by $i(a, b)$. A *multicurve* on S is a disjoint union of finitely many curves on S , and it is *essential* if all its components are.

By “homeomorphism of S ” we always mean an orientation-preserving homeomorphism $S \rightarrow S$. The *mapping class group* of S , denoted by $\text{Mcg}(S)$, is the group of homeomorphisms of S modulo isotopies.

Complexity of surfaces. In this article, we aim to provide algorithms whose running times depend polynomially on the complexity of a surface. There are two essentially equivalent measures of complexity of a surface S : its Euler characteristic $\chi = 2 - 2g - p$, and the integer

$$\xi = 3g - 3 + p,$$

which counts, for instance, the maximum number of disjoint non-parallel essential curves on S (assuming that $\xi \geq 1$). For the sake of consistency, we will employ ξ as our notion of complexity of S , although we remark that, for sufficiently complex surfaces, the two integers ξ and $-\chi$ are equivalent up to a multiplicative constant.

For surfaces with $\xi \leq 1$ – namely, the sphere with at most four punctures and the torus with at most one puncture – the problems we are trying to solve are already well-understood. In fact, for these surfaces, the mapping class group and the curve graph are either essentially trivial (for the sphere with at most three punctures), or easily describable in terms of $\mathrm{SL}(2, \mathbb{Z})$ and the Farey graph (see [12, §2.2] and [3]). Therefore, unless otherwise stated, we will work with the underlying assumption that our ambient surface (always denoted by S) has complexity $\xi \geq 2$.

Triangulations. To provide a surface as input to an algorithm, it is necessary to represent it combinatorially, and triangulations frequently offer the most suitable approach for this purpose. We begin by describing the pertinent concepts for punctured surfaces, while deferring the discussion of closed surfaces to a subsequent paragraph.

Let us think of the punctures of a punctured surface S as marked points on a closed surface. An *ideal triangulation* (or simply a *triangulation*) of S is an ordered set $\mathcal{T} = \{e_1, \dots, e_n\}$ where each e_i is an arc embedded in S , such that

- i the endpoints of each arc are punctures,
- ii the interiors of any two arcs are disjoint, and
- iii each component Z of $S \setminus (e_1 \cup \dots \cup e_n)$ is a *triangle*, in the sense that there is a map $\Delta \rightarrow S$ from a Euclidean 2-simplex Δ to S which sends each vertex of Δ to a puncture of S , the interior of each edge of Δ homeomorphically to the interior of an arc of \mathcal{T} , and the interior of Δ homeomorphically to Z , and moreover Z does not contain any punctures.

The arcs e_1, \dots, e_n are called *edges* of \mathcal{T} ; an easy counting argument shows that there are exactly $n = -3\chi$ edges in any ideal triangulation of S . We will denote by $|\mathcal{T}|$ the number of edges of \mathcal{T} .

Note that a triangulation can be described combinatorially by providing a set of punctures, an ordered list of edges – along with the information of which punctures are the endpoints of each edge – and a set of triangles – along with the information of the sequence of edges forming the boundary of each triangle. We will always assume that surfaces are given as inputs to our algorithms in this form.

When S is a closed surface, we can introduce a fictitious marked point and consider triangulations having this point as their only vertex. We call them *one-vertex triangulations* (or simply *triangulations*) of S ; this is the combinatorial representation of closed surfaces that our algorithms will accept as input.

Normal curves. Continuing our endeavour to establish combinatorial representations of the topological objects we wish to manipulate, we now shift our focus to curves on the surface S .

Let us fix a triangulation \mathcal{T} of S , one-vertex or ideal depending on whether S is closed or not. If a is an essential curve (or, in fact, an essential multicurve) on S , then it can be isotoped so that its intersection with each triangle Δ of \mathcal{T} is a collection of arcs, each of which connects distinct edges of Δ ; we call such a (multi)curve *normal*. The typical intersection of a normal curve with a triangle of \mathcal{T} is illustrated in Figure 2.1.

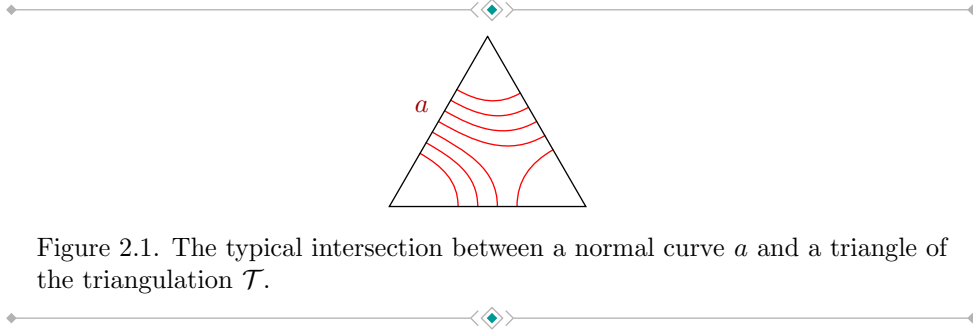


Figure 2.1. The typical intersection between a normal curve a and a triangle of the triangulation \mathcal{T} .

A normal curve a can be efficiently encoded by recording how many times it intersects each edge e of \mathcal{T} ; we denote this number by $\mathcal{T}(a)_e$. Recalling that edges of \mathcal{T} are ordered, we can collect the non-negative integers $\{\mathcal{T}(a)_e : e \in \mathcal{T}\}$ into a vector $\mathcal{T}(a) \in \mathbb{Z}_{\geq 0}^n$, where $n = |\mathcal{T}|$ is the number of edges of \mathcal{T} .

This representation is unambiguous, in the sense that we can recover the curve a up to isotopy from the vector $\mathcal{T}(a)$. It is also efficient, in the sense that it can be used to represent exponentially complicated curves with a linear amount of information. More precisely, we define the *complexity* of a normal curve a with respect to the triangulation \mathcal{T} to be the number

$$\|a\|_{\mathcal{T}} = \sum_{e \in \mathcal{T}} \log(\mathcal{T}(a)_e + 1),$$

where here – and everywhere else – we take the base of the logarithm to be 2. We will omit the subscript \mathcal{T} if the triangulation can be inferred from the context.

Note that the complexity of a is roughly proportional to the number of digits that are required to store the vector $\mathcal{T}(a)$, but is clearly logarithmic in the number of intersections between a and the edges of \mathcal{T} . When we say that an algorithm takes as input an essential curve a on a surface S , we will actually mean that it takes the vector $\mathcal{T}(a)$, where \mathcal{T} is the triangulation representing S . Likewise, all actions performed by our algorithms on essential curves will only operate on this vector representation.

Sometimes, it is convenient to use an alternative notion of complexity of normal curves, defined by

$$\|a\|_{\mathcal{T},1} = \log\left(\sum_{e \in \mathcal{T}} \mathcal{T}(a)_e\right)$$

for a normal curve a ; like above, we will omit the subscript \mathcal{T} unless needed. Note that the two complexities are related by the inequalities

$$\|a\|_1 \leq \|a\| \leq |\mathcal{T}| \cdot \|a\|_1.$$

We conclude by remarking that, when S is punctured, the vector $\mathcal{T}(a)$ is invariant under isotopy of a , in the sense that if a and b are normal curves on S and are isotopic, then $\mathcal{T}(a) = \mathcal{T}(b)$. This property immediately gives a linear-time algorithm in $\|a\| + \|b\|$ to decide whether two curves a and b are isotopic. The situation is not as straightforward

when S is closed – essentially because we allow isotopies across the vertex of \mathcal{T} – but a result of Lackenby [22, Theorem 1.2] gives us the following proposition.

Proposition 2.1. *Let a and b be essential curves on a surface S . There is an algorithm to decide if a and b are isotopic. The running time of the algorithm is polynomial in ξ , $\|a\|$, and $\|b\|$.*

Flip and twist graph. A convenient framework to describe the relations between different triangulations of a surface is given by the flip and twist graph, introduced for algorithmic purposes by Bell [6]. Suppose first that the surface S is punctured. The *flip and twist graph* of S is a graph $\mathcal{G}(S)$ whose vertices are isotopy classes of triangulations of S , and edges connect triangulations which are related by certain elementary moves, namely reorderings, flips, and powers of Dehn twists. Roughly speaking, the length of an edge is proportional to the computational complexity of performing the corresponding move. More precisely, we can describe the three types of edges in $\mathcal{G}(S)$ as follows.

- If two triangulations \mathcal{T} and \mathcal{T}' have the same edges up to reordering, they are joined by an edge of length 1 in $\mathcal{G}(S)$.
- If e is an edge of a triangulation \mathcal{T} which is adjacent to two different triangles Δ_1 and Δ_2 , a *flip* about e yields a new triangulation \mathcal{T}' obtained by replacing e with the opposite diagonal of the square $\Delta_1 \cup \Delta_2$ (see [4, §2.2]). In the graph $\mathcal{G}(S)$, the triangulations \mathcal{T} and \mathcal{T}' are joined by an edge of length 1.
- Given an essential curve a on S , denote by T_a the Dehn twist of S about a . If \mathcal{T} is a triangulation of S and \mathcal{T}' is the image of \mathcal{T} under T_a^k for some integer $k \neq 0$, then there is an edge of length $\|a\|_{\mathcal{T},1} + \log |k|$ in $\mathcal{G}(S)$ between \mathcal{T} and \mathcal{T}' .

When S is closed, its flip and twist graph is the flip and twist graph of the surface obtained by adding a fictitious marked point to S .

The graph $\mathcal{G}(S)$ is connected, as proved by Hatcher [19], even if we ignore the edges corresponding to powers of Dehn twists. If we have a path of length ℓ in $\mathcal{G}(S)$ between two triangulations \mathcal{T} and \mathcal{T}' , then we can perform a change of coordinates of normal curves between the two triangulations in polynomial time in ℓ . More precisely, if a is a normal curve on S with respect to the triangulation \mathcal{T} , then we can compute the vector $\mathcal{T}'(a)$ in polynomial time in ξ , ℓ , and $\|a\|_{\mathcal{T}}$. This follows from [4, Proposition 2.2.1] for flips and [29, Theorem 4.1] for powers of Dehn twists.

Representing homeomorphisms. There are a few reasonable combinatorial representations of homeomorphisms $f: S \rightarrow S$ that are amenable to algorithmic manipulations. The algorithms we describe in Section 4 are largely agnostic when it comes to this choice; our only requirement is that the selected representation comes equipped with a notion of complexity $\|-\|$ satisfying the following properties.

- i** Given a homeomorphism $f: S \rightarrow S$ of complexity $\|f\|$ and an essential curve a on S , there is an algorithm to compute the curve $f(a)$ in polynomial time in ξ , $\|f\|$, and $\|a\|$.
- ii** For a homeomorphism $f: S \rightarrow S$ and an essential curve a on S , the complexities of f , a , and $f(a)$ are related by the inequality

$$\|f(a)\|_1 \leq \text{poly}(\xi) \cdot \text{poly}(\|f\|) + \|a\|_1.$$

Note that, in Condition **ii**, we are using the complexity $\|-\|_1$ of curves.

In order to provide a concrete example, we now describe how homeomorphisms can be represented as paths in the flip and twist graph of S . Let \mathcal{T} be a fixed triangulation of the surface S . If $f: S \rightarrow S$ is a homeomorphism, then $f(\mathcal{T})$ is also a triangulation of S ; if S is closed, up to isotopy we may assume that f fixes the vertex of \mathcal{T} . Since the flip and twist graph $\mathcal{G}(S)$ is connected, there is a path in $\mathcal{G}(S)$ from \mathcal{T} to $f(\mathcal{T})$, and we can use this path as a combinatorial representation of the homeomorphism f . We then define the complexity $\|f\|$ to be the length of this path.

We have already discussed how, given an essential curve a on S , we can compute the curve $f(a)$ in polynomial time in ξ , $\|f\|$, and $\|a\|$. As far as Condition [ii](#) is concerned, it is not hard to check that the inequality

$$\|f(a)\|_1 \leq 3\|f\| + \|a\|_1$$

is satisfied for every essential curve a on S .

2.2 Curve graph

One of the main goals of this article is to give an efficient algorithm to coarsely compute distances in the curve graph of a surface, which we now define following Harvey [18]. Let S be a surface with $\xi(S) \geq 2$. The *curve graph* of S is a graph whose vertex set is the set $\mathcal{C}_0(S)$ of isotopy classes of essential curves in S . Two classes of curves are connected by an edge if they admit disjoint representatives. By stipulating that every edge has length 1, the curve graph becomes a metric space, and we can talk about the distance $\text{dist}_{\mathcal{C}(S)}(a, b)$ between two (isotopy classes of) curves a and b .

The curve graph is connected, and in fact we can give a more precise upper bound for the distance between two curves in terms of their intersection number. The following proposition combines several known upper bounds of this form.

Proposition 2.2. *Let $a, b \in \mathcal{C}_0(S)$ be essential curves with $i(a, b) > 0$. Then*

$$\text{dist}_{\mathcal{C}(S)}(a, b) \leq \min \left\{ i(a, b) + 1, 2 \log i(a, b) + 2, \frac{6}{\log \xi} \cdot \log i(a, b) + 2 \right\}.$$

If $i(a, b) \leq 2$ then we have the better bound

$$\text{dist}_{\mathcal{C}(S)}(a, b) \leq 2.$$

Proof. Firstly, note that if a and b intersect at most twice, then they cannot fill a surface of complexity $\xi \geq 2$, and hence $\text{dist}_{\mathcal{C}(S)}(a, b) \leq 2$. In the general case, the inequality $\text{dist}_{\mathcal{C}(S)}(a, b) \leq i(a, b) + 1$ is proved in [9, Lemma 1.1]. The logarithmic bound $\text{dist}_{\mathcal{C}(S)}(a, b) \leq 2 \log i(a, b) + 2$ is known as Hempel's lemma [20, Lemma 2.1], and is proved in [30, Lemma 1.21] in the punctured case.

Finally, the last bound is a consequence of Lemma 2.1 of [10]. In fact, it follows from Bowditch's result that

$$\text{dist}_{\mathcal{C}(S)}(a, b) \leq \frac{2}{\log(2g + p - 3) - 1} \cdot \log i(a, b) + 2;$$

we get the desired inequality by noting that $\log \xi \leq 3 \max\{1, \log(2g + p - 3) - 1\}$. \square

When S is the sphere with 4 punctures or the torus with 0 or 1 punctures, the definition above would yield a totally disconnected graph. Therefore, in these cases, we

change the definition by stipulating that two classes of curves are connected by an edge if their intersection number is 1 – in the case of the torus with at most one puncture – or 2 – in the case of the sphere with 4 punctures.

When S is an annulus (that is, a twice-punctured sphere), we adopt yet another definition. In the scope of this article, we will only be concerned with curve graphs of annuli when dealing with the subsurface projection map; therefore, we defer the definition of $\mathcal{C}(S)$ in the annular case to Section 4.2.

Finally, let us introduce the following notation: if A and B are non-empty subsets of $\mathcal{C}_0(S)$, we define

$$\begin{aligned} \text{dist}_{\mathcal{C}(S)}(A, B) &= \min\{\text{dist}_{\mathcal{C}(S)}(a, b) : a \in A, b \in B\}, \\ \text{diam}_{\mathcal{C}(S)}(A) &= \sup\{\text{dist}_{\mathcal{C}(S)}(a, a') : a, a' \in A\}, \\ \text{diam}_{\mathcal{C}(S)}(A, B) &= \sup\{\text{dist}_{\mathcal{C}(S)}(a, b) : a \in A, b \in B\}. \end{aligned}$$

2.3 Combinatorics of train tracks

Train tracks are one of the main tools to compute lower bounds for distances in the curve graph. We give a concise review of the pertinent definitions and properties, and then introduce new notations tailored to our algorithmic purposes. We will try to clearly differentiate between concepts derived from established literature and novel definitions. Our primary reference for train tracks, and in particular for their role in the computation of distances in the curve graph, is [24, §4].

Definition. A *train track* on a surface S is an embedded 1-complex $\tau \subseteq S$ such that every edge (called a *branch*) is smooth and has well-defined tangent vectors at its endpoints; at each vertex (called a *switch*) the incident branches are mutually tangent, and there is at least one branch for each of the two possible tangent directions. The valence of each switch is at least 3, except for exactly one bivalent switch on each simple closed curve component of τ . Finally, the connected components of $S \setminus \tau$ (which we call *complementary regions*) are surfaces whose boundaries consist of smooth arcs interleaved with cusps; we require that no complementary region is a nullgon (that is, a smooth disc), a monogon (that is, a disc with one cusp), a bigon (that is, a disc with two cusps), a once-punctured nullgon, or a smooth annulus.

Directions and ends. We denote by $\mathcal{S}(\tau)$ and $\mathcal{B}(\tau)$ the sets of switches and branches of τ respectively. A simple counting argument (see [28, Corollary 1.1.3]) shows that

$$|\mathcal{S}(\tau)| \leq 4\xi \quad \text{and} \quad |\mathcal{B}(\tau)| \leq 6\xi.$$

A *direction* of a switch $s \in \mathcal{S}(\tau)$ is a non-zero tangent vector to τ at s , defined up to multiplication by a positive scalar; in particular, each switch has two directions. If η is a direction of s , we denote the opposite direction by $\bar{\eta}$.

We say that a branch of τ is *returning* if it is adjacent to a single switch s , and its two endpoints have opposite directions at s .

A *half-branch* of a branch $b \in \mathcal{B}(\tau)$ is a component of $b \setminus q$ for a point q in the interior of b . An *end* of b is an equivalence class of half-branches of b , where two half-branches are equivalent if their intersection is still a half-branch of b . It is clear that every branch b has two ends; we will sometimes call one of them b_\bullet , and implicitly refer to the other one as b^\bullet – or vice versa.

The switch *incident* to an end e , denoted by $\text{switch}(e)$, is the unique switch of τ lying in the closure of a half-branch representing e . The *direction* of e , denoted by $\text{dir}(e)$, is the

direction of $\text{switch}(e)$ represented by the tangent vector at the endpoint of a half-branch representing e .

Denote the set of ends of τ – that is, the set of ends of branches of τ – by $\mathcal{E}(\tau)$. For every switch s of τ and every direction η of s , we define

$$\mathcal{E}_\eta(\tau) = \{e \in \mathcal{E}(\tau) : \text{dir}(e) = \eta\}.$$

◇ Subtracks and carrying. For our purposes, a *carrying map* between two train tracks (which may be multicurves) τ' and τ is a map $\varphi: S \rightarrow S$ homotopic to the identity such that:

- i** $\varphi(\tau') \subseteq \tau$;
- ii** $\varphi(s') \in \mathcal{S}(\tau)$ for every $s' \in \mathcal{S}(\tau')$;¹
- iii** the differential of the restriction $\varphi|_{\tau'}: \tau' \rightarrow \tau$ is nowhere vanishing.

If there exists a carrying map between τ' and τ , we say that τ' is *carried* by τ , and we write $\tau' \prec \tau$. A very special case of carrying occurs when τ' is a subset of τ ; in this case, we say that τ' is a *subtrack* of τ , and we write $\tau' < \tau$. Whenever we write $\tau' \prec \tau$, we always implicitly assume that the carrying is induced by a specific carrying map, which we denote by $\varphi_{\tau' \prec \tau}$.

Our preferred way of depicting a carrying $\tau' \prec \tau$ is shown in Figure 2.2: we draw a thickened neighbourhood of τ instead of τ itself, and we represent τ' so that it is contained in this neighbourhood.




Figure 2.2. A representation of a carrying $\tau' \prec \tau$. We draw τ' inside a thickened neighbourhood of τ .

◇ Pre-measures. A *pre-measure* on τ is a function

$$\mu: \mathcal{B}(\tau) \longrightarrow \mathbb{Z}_{\geq 0},$$

where $\mathbb{Z}_{\geq 0}$ denotes the set of non-negative integers. We say that μ *fills* τ if $\mu(b) > 0$ for every $b \in \mathcal{B}(\tau)$. With slight abuse of notation, if e is an end of a branch b , we will write $\mu(e)$ for $\mu(b)$.

For every branch b of τ , define $\mathbb{1}_b$ to be the pre-measure assigning weight 1 to b and weight 0 to every other branch of τ . Moreover, if μ and μ' are pre-measures on τ , we will write $\mu + \mu'$ for the pre-measure obtained by adding μ and μ' pointwise, and $\mu \geq \mu'$ as a shorthand for “ $\mu(b) \geq \mu'(b)$ for every $b \in \mathcal{B}(\tau)$ ”.

¹This condition is not part of the usual definition of carrying map in the literature.

Two remarks are in order. Firstly, the notion of pre-measure is not found in the literature. In fact, the only reason why we give this definition instead of directly jumping to that of measure (see below) is that it is sometimes convenient in the description of our algorithms to refer to the pre-measures $\mathbb{1}_b$. Secondly, it is common in the literature to define measures as taking non-negative real values. However, since we will focus solely on curves (and not on laminations), we have chosen a more combinatorial approach by restricting our definition to integral (pre-)measures.

◇ Measures. A pre-measure μ on τ is a *measure* if for every switch $s \in \mathcal{S}(\tau)$ and every direction η of τ the equality

$$\sum_{e \in \mathcal{E}_\eta(\tau)} \mu(e) = \sum_{e \in \mathcal{E}_{\bar{\eta}}(\tau)} \mu(e)$$

holds. For a measure μ and a switch s , we denote the quantity defined by either side of the above equation by $\mu(s)$.

Every non-zero measure μ on τ represents a non-empty essential multicurve carried by τ . This multicurve is constructed by taking $\mu(b)$ parallel copies of b for each branch $b \in \mathcal{B}(\tau)$, and gluing corresponding arcs at switches. Every essential multicurve carried by τ is represented by some measure on τ , and different measures always represent non-isotopic multicurves (see, for instance, Proposition 2.7.4 in Penner and Harer [28]’s book). For this reason, we will often blur the distinction between measures on a train track and essential multicurves carried by it.

We call a pair (τ, μ) a *measured train track* if τ is a train track and μ is a measure on τ ; we say that (τ, μ) is *connectedly fully measured* if μ fills τ and represents a connected curve. A customary way to graphically represent a measured train track (τ, μ) is to draw a thickened neighbourhood of τ instead of τ itself, where the thickness of a branch is proportional to the value assigned to it by μ .

For a train track τ , we define $P(\tau)$ to be the set of measures on τ . Moreover, we let

$$\text{int } P(\tau) = \{\mu \in P(\tau) : \mu \text{ fills } \tau\}.$$

A train track τ is *recurrent* if $\text{int } P(\tau)$ is non-empty.

Given two train tracks $\tau' \prec \tau$, we say that τ' *fills* τ if $\text{int } P(\tau') \subseteq \text{int } P(\tau)$; note that talking about containment of these two sets makes sense if we interpret them as sets of isotopy classes of multicurves on S .

◇ Fundamental curves. A *fundamental measure*² on a train track τ is a non-zero measure $\nu \in P(\tau)$ that cannot be written as the sum of two non-zero measures on τ . We denote the set of fundamental measures on τ by $\mathcal{F}(\tau)$. A fundamental measure will obviously represent a connected curve on S ; therefore, we will sometimes refer to fundamental measures as “fundamental curves”.

It is easy to see that, if ν is a fundamental measure on τ , then $\nu(s) \leq 2$ for every switch $s \in \mathcal{S}(\tau)$. In fact, assume for contradiction that $\nu(s) \geq 3$ for some switch s of τ , and fix an orientation for the curve a represented by ν . The curve a will run through s at least twice with the same orientation, as depicted in Figure 2.3a. With a local cut-and-paste operation near s , as shown in Figure 2.3b, we transform a into two immersed curves a_1 and a_2 ; these curves may have self-intersections, but they give rise to two non-zero measures ν_1 and ν_2 by counting how many times they run along branches of τ . It is clear that $\nu = \nu_1 + \nu_2$, hence ν is not fundamental.

²Note that this definition is less restrictive than that of *vertex measure* usually found in the literature; for this reason, we have decided to employ a different name.

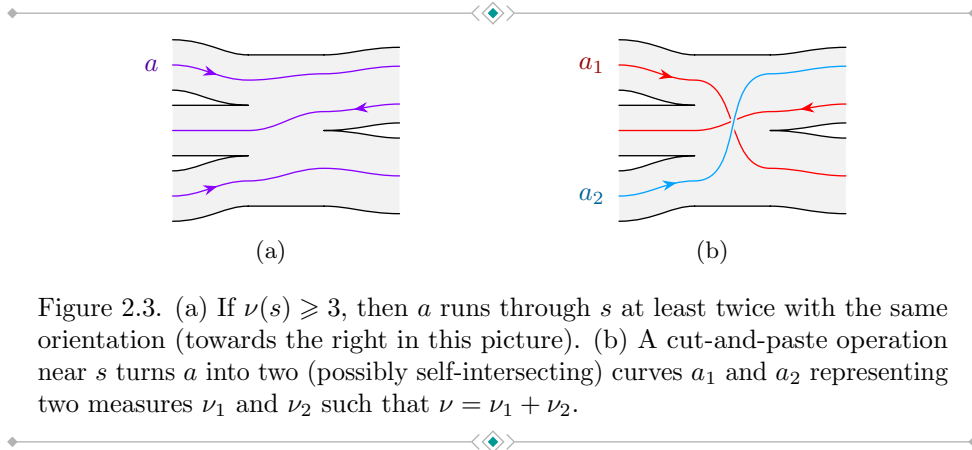


Figure 2.3. (a) If $\nu(s) \geq 3$, then a runs through s at least twice with the same orientation (towards the right in this picture). (b) A cut-and-paste operation near s turns a into two (possibly self-intersecting) curves a_1 and a_2 representing two measures ν_1 and ν_2 such that $\nu = \nu_1 + \nu_2$.

Largeness and diagonal extensions. A train track τ is *large* if its complementary regions are topological discs with at most one puncture; in other words, if every essential curve in S intersects τ . Note that if τ is large and carried by a train track τ' , then τ' is large as well.

We say that a train track ρ is a *diagonal extension* of a train track τ if $\tau < \rho$ and $\rho \setminus \tau$ consists of branches whose ends terminate in cusps of complementary regions of τ . If τ' is a large recurrent train track which is carried by τ and fills τ , then every diagonal extension of τ' is carried by a diagonal extension of τ (see the proof of Lemma 4.2 in [24]).

For a large train track τ , we define the following sets of measures/multicurves:

$$\begin{aligned} \text{PE}(\tau) &= \bigcup \{P(\rho) : \rho \text{ is a recurrent diagonal extension of } \tau\}, \\ \text{int PE}(\tau) &= \{\mu \in \text{PE}(\tau) : \mu(b) > 0 \text{ for every } b \in \mathcal{B}(\tau)\}, \\ \text{PN}(\tau) &= \bigcup \{\text{PE}(\sigma) : \sigma \text{ is a large recurrent subtrack of } \tau\}, \\ \text{int PN}(\tau) &= \bigcup \{\text{int PE}(\sigma) : \sigma \text{ is a large recurrent subtrack of } \tau\}. \end{aligned}$$

Distances in $\mathcal{C}(S)$. Let us begin with a trivial observation: any two curves carried by a non-large train track lie at distance at most two in $\mathcal{C}(S)$. This follows from the fact that a train track τ is not large precisely when there is an essential curve on S disjoint from τ .

Next, let us give an upper bound for the intersection number of curves carried by a train track τ . If μ and μ' are measures on τ representing curves a and a' respectively, then

$$i(a, a') \leq \sum_{s \in \mathcal{S}(\tau)} \mu(s) \cdot \mu'(s).$$

In fact, a and a' can be isotoped so that they run parallel to the branches of τ , and only intersect near the switches; in particular, near a switch $s \in \mathcal{S}(\tau)$, they intersect at most $\mu(s) \cdot \mu'(s)$ times, as depicted in Figure 2.4.

We now state what is probably the most crucial result on which our algorithm for computing coarse distances in $\mathcal{C}(S)$ relies.




Figure 2.4. The curves a and a' can be realised so that they run parallel to each other along branches of τ , and intersect at most $\mu(s) \cdot \mu'(s)$ near a switch s .

Proposition 2.3. *Let τ be a large recurrent train track, and let $a \in \text{int PN}(\tau)$ be an essential curve. Let $b \in \mathcal{C}_0(S)$ be any curve such that $i(a, b) = 0$. Then $b \in \text{PN}(\tau)$.*

This proposition first appeared as Lemma 4.4 in [24], with the additional constraint that τ should be *birecurrent*. However, the assumption of *transverse recurrence* can be dropped, as shown by Gadre and Tsai [14, Lemma 3.2]; they attribute the proof to Leininger. Note that Gadre and Tsai state the lemma for int PE and PE instead of int PN and PN . However, our Proposition 2.3 is an immediate consequence of their result. In fact, $a \in \text{int PN}(\tau)$ means that $a \in \text{int PE}(\sigma)$ for a large recurrent subtrack $\sigma < \tau$. By Lemma 3.2 of [14], this implies that $b \in \text{PE}(\sigma)$, and by definition $\text{PE}(\sigma) \subseteq \text{PN}(\tau)$.

An easy consequence of Proposition 2.3 allows us to provide lower bounds for distances in the curve graph in terms of specific nested sequences of train tracks.

Corollary 2.4. *Let $\tau_0 \succ \tau_1 \succ \dots \succ \tau_n$ be a sequence of large recurrent train tracks such that*

$$\text{PN}(\tau_i) \subseteq \text{int PN}(\tau_{i-1}) \quad \text{for every } 1 \leq i \leq n.$$

Then there is a curve $a \in \mathcal{F}(\tau_0)$ such that

$$\text{dist}_{\mathcal{C}(S)}(a, \mathcal{F}(\tau_n)) \geq n.$$

Proof. Let b be a fundamental curve of τ_n ; in particular, note that $b \in \text{PN}(\tau_n)$. Making use of Proposition 2.3, we can prove by induction on $0 \leq k < n$ that every essential curve whose distance from b is at most k lies in $\text{int PN}(\tau_{n-k-1})$. Let us now fix a curve a representing a *vertex measure* of τ_0 according to the definition in [24, §4.1]. Note that a is a fundamental curve of τ_0 . By Lemma 3.5 in [26], the curve a does not lie in $\text{int PN}(\tau_0)$; by the above discussion, this implies that $\text{dist}_{\mathcal{C}(S)}(a, b) \geq n$. Since a can be chosen independently of b , this concludes the proof. \square

Co-orientations and height. Let us remark that all the terminology we will introduce from now on is not standard in the literature; as anticipated, some of these definitions are specifically tailored to our algorithms, and would probably not be useful in most other contexts.

Loosely speaking, a co-orientation of a switch $s \in \mathcal{S}(\tau)$ is a unit tangent vector orthogonal to τ at s , as depicted in Figure 2.5a. More formally, a *co-orientation* of s is a tangent vector to S at the point s which is not tangent to τ , defined up to the following equivalence relation: two co-orientations ω and ω' are the same if, for every direction η of s , the bases $\{\eta, \omega\}$ and $\{\eta, \omega'\}$ of the tangent space to S at s are related by an

orientation-preserving linear isomorphism. Clearly, every switch has two co-orientations; if ω is a co-orientation of s , we denote the opposite co-orientation by $\bar{\omega}$.

Fix a switch s and a direction η of s . A co-orientation ω of s defines a linear order on the set of ends $\mathcal{E}_\eta(\tau)$, given by enumerating the ends following the vector ω ; see Figure 2.5b for an example. For two ends $e_1, e_2 \in \mathcal{E}(\tau)$, we will say that e_1 is ω -smaller than e_2 (and write $e_1 <_\omega e_2$) if $\text{switch}(e_1) = \text{switch}(e_2) = s$, $\text{dir}(e_1) = \text{dir}(e_2)$, and e_1 is smaller than e_2 according to the order defined above.




Figure 2.5. (a) A co-orientation ω of a switch s can be informally thought of as a unit tangent vector orthogonal to τ at s . (b) The linear order induced on $\mathcal{E}_\eta(\tau)$ by the co-orientation ω .

We denote by $\mathcal{E}_\eta(\tau)^\omega$ the set $\mathcal{E}_\eta(\tau)$ endowed with the linear order $<_\omega$. For $1 \leq i \leq |\mathcal{E}_\eta(\tau)|$, we denote by $\mathcal{E}_\eta(\tau)_i^\omega$ the i -th element of $\mathcal{E}_\eta(\tau)^\omega$. We will say that two ends of τ are *consecutive* if they are of the form $\mathcal{E}_\eta(\tau)_i^\omega$ and $\mathcal{E}_\eta(\tau)_{i+1}^\omega$ for some switch $s \in \mathcal{S}(\tau)$, some direction η of s , some co-orientation ω of s , and some integer i ; note that this condition does not depend on the choice of ω .

If μ is a pre-measure on τ , s a switch of τ , ω a co-orientation of s , and e an end with $\text{switch}(e) = s$, then we define the *height* of e as

$$\text{height}_\omega(e) = \sum_{\substack{e' \in \mathcal{E}(\tau) \\ e' >_\omega e}} \mu(e').$$

The pre-measure μ is omitted from the notation, but will always be clear from the context. The name “height” suggests the convention that, in our pictures, co-orientations will always point downwards.

Carrying revisited. Recall that, according to our definition, a carrying map φ between two train tracks τ' and τ sends switches to switches. As a consequence, there are induced maps – which we keep calling φ – sending directions, co-orientations, ends, and pre-measures on τ' to directions, co-orientations, ends, and pre-measures on τ . Let us now state two properties of these maps.

- If s' is a switch of τ' , η' a direction of s' , ω' a co-orientation of s' , $\eta = \varphi(\eta')$, and $\omega = \varphi(\omega')$, then the map

$$\varphi: \mathcal{E}_{\eta'}(\tau')^{\omega'} \longrightarrow \mathcal{E}_\eta(\tau)^\omega$$

is order-preserving. This is essentially obvious, as the reader can see from Figure 2.6.

- The map φ sending pre-measures on τ' to pre-measures on τ is \mathbb{Z} -linear, and is uniquely determined by the pre-measures $\{\varphi(\mathbb{1}_{b'}) : b' \in \mathcal{B}(\tau')\}$. As a consequence, it sends measures to measures, and there is an induced \mathbb{Z} -linear map

$$\varphi: P(\tau') \longrightarrow P(\tau).$$




Figure 2.6. The function $\mathcal{E}_{\eta'}(\tau')^{\omega'} \rightarrow \mathcal{E}_{\eta}(\tau)^{\omega}$ induced by a carrying map φ is order-preserving.

In the context of a carrying $\tau' \prec \tau$, we will say that a branch $b \in \mathcal{B}(\tau)$ is *untouched* in τ' if $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'}) = \mathbb{1}_b$ for some branch $b' \in \mathcal{B}(\tau')$. We will say that an end $e \in \mathcal{E}(\tau)$ is *persistent* in τ' if $\varphi_{\tau' \prec \tau}(e') = e$ for some end $e' \in \mathcal{E}(\tau')$.

If (τ, μ) and (τ', μ') are measured train tracks, we write $(\tau', \mu') \prec (\tau, \mu)$ to signify that $\tau' \prec \tau$ and $\varphi_{\tau' \prec \tau}(\mu') = \mu$.

One-switch train tracks. A *one-switch train track* is a train track with only one switch. One-switch train tracks will be particularly relevant for our algorithms, since they are more convenient to work with than general train tracks. The main reason for this is that they only have $\mathcal{O}(\xi^2)$ fundamental curves, which can be described explicitly. In fact, if τ is a one-switch train track, then

$$\begin{aligned} \mathcal{F}(\tau) = & \{ \mathbb{1}_b : b \in \mathcal{B}(\tau) \text{ is returning} \} \\ & \cup \{ \mathbb{1}_{b_1} + \mathbb{1}_{b_2} : b_1, b_2 \in \mathcal{B}(\tau) \text{ are not returning and } \text{dir}(b_1^\bullet) \neq \text{dir}(b_2^\bullet) \}. \end{aligned}$$

Figure 2.7 shows the two types of fundamental curves of a one-switch train track.




Figure 2.7. The two types of fundamental curves of a one-switch train track: the curve a is represented by $\mathbb{1}_b$ for a returning branch b , while c is represented by $\mathbb{1}_{b_1} + \mathbb{1}_{b_2}$ for non-returning branches b_1 and b_2 .

We observe that two fundamental curves of a one-switch train track τ intersect at most 4 times. The bound in Proposition 2.2 implies that

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau)) \leq 5. \quad (1)$$

Deep nesting. Let $\tau' \prec \tau$ be a pair of one-switch train tracks. We say that τ' is *deeply nested* in τ if exactly two ends of τ are persistent in τ' . Figure 2.8a shows what two deeply nested train tracks typically look like near their unique switches. In particular, the train track τ has two branches b_1 and b_2 (not necessarily distinct) with ends b_1^\bullet and b_2^\bullet satisfying $\text{dir}(b_1^\bullet) \neq \text{dir}(b_2^\bullet)$ and $\varphi_{\tau' \prec \tau}(e') \in \{b_1^\bullet, b_2^\bullet\}$ for every end e' of τ' .

We now state two properties of deeply nested train tracks which will be useful in the proof of Theorem 3.7. We use the notation we introduced in the previous paragraph.

Lemma 2.5. *Let $\tau' \prec \tau$ be deeply nested one-switch train tracks.*

- 1 *There is a fundamental measure ν on τ , supported on the branches b_1 and b_2 , such that $\varphi_{\tau' \prec \tau}(\mu') \geq \nu$ for every non-zero measure μ' on τ' .*
- 2 *Suppose that τ' is large, recurrent, and fills τ . Then every recurrent diagonal extension ρ' of τ' is carried by and deeply nested in a recurrent diagonal extension ρ of τ ; the two ends of ρ which are persistent in ρ' are b_1^\bullet and b_2^\bullet .*

Proof. We prove the two statements separately.

- 1 Let μ' be a non-zero measure on τ' . There must be two ends e'_1 and e'_2 of τ' with $\text{dir}(e'_1) \neq \text{dir}(e'_2)$ such that $\mu'(e'_1) > 0$ and $\mu'(e'_2) > 0$. Without loss of generality, assume that $\varphi_{\tau' \prec \tau}(e'_1) = b_1^\bullet$ and $\varphi_{\tau' \prec \tau}(e'_2) = b_2^\bullet$. If we let $\mu = \varphi_{\tau' \prec \tau}(\mu')$, it immediately follows that $\mu(b_1) > 0$ and $\mu(b_2) > 0$. If b_1 (respectively b_2) is returning, we can take $\nu = \mathbb{1}_{b_1}$ (respectively $\mathbb{1}_{b_2}$). Otherwise, we take $\nu = \mathbb{1}_{b_1} + \mathbb{1}_{b_2}$. Either way, we have that ν is a fundamental measure on τ and $\mu \geq \nu$.
- 2 We already know, from the proof of [24, Lemma 4.2], that ρ' is carried by a recurrent diagonal extension ρ of τ . Now, let b' be a branch in $\mathcal{B}(\rho') \setminus \mathcal{B}(\tau')$, and let b'_\bullet be one of its ends. Since ρ' is a diagonal extension of τ' , there are two consecutive ends e'_1 and e'_2 of τ' such that $e'_1 <_{\omega'} b'_\bullet <_{\omega'} e'_2$ for some orientation ω' of the unique switch of τ' ; in other words, b'_\bullet is contained in the cusp of $S \setminus \tau'$ defined by the two ends e'_1 and e'_2 , as shown in Figure 2.8b. But $\varphi_{\rho' \prec \rho}(e'_1) = \varphi_{\rho' \prec \rho}(e'_2) = b_i^\bullet$ for some $i \in \{1, 2\}$, and $\varphi_{\rho' \prec \rho}$ preserves the order of ends, so $\varphi_{\rho' \prec \rho}(b'_\bullet) = b_i^\bullet$ as well. It follows that ρ' is deeply nested in ρ , and the only two persistent ends of ρ in ρ' are b_1^\bullet and b_2^\bullet . \square

While the significance of the “deeply nested” condition will become apparent in the proof of Theorem 3.7, it is easy to get lost in the technical details of the argument, and the reader might be left wondering why we have introduced this concept in the first place. At a very high level, the reason is exemplified by the following situation.

Let $\tau'' \prec \tau$ be large recurrent one-switch train tracks. In light of Corollary 2.4, we might be interested in deciding whether the condition $\text{PN}(\tau'') \subseteq \text{int PN}(\tau)$ holds. This containment is challenging to check in polynomial time in ξ , mainly because the number of diagonal extensions of large subtracks of τ'' could a priori be doubly exponential in ξ . Suppose, however, that there is a train track τ' with $\tau'' \prec \tau' \prec \tau$ such that τ'' is deeply nested in τ' ; let b'_1 and b'_2 be the two (not necessarily distinct) branches of τ' having an end which is persistent in τ'' . The content of Lemma 2.5 is essentially the following:




Figure 2.8. (a) Two deeply nested train tracks $\tau' \prec \tau$; we have that $\varphi_{\tau' \prec \tau}(e') \in \{b_1^\bullet, b_2^\bullet\}$ for every end e' of τ' . (b) If the end b_\bullet^\bullet is contained in the cusp of τ' defined by e'_1 and e'_2 , and $\varphi_{\tau' \prec \tau}(e'_1) = \varphi_{\tau' \prec \tau}(e'_2) = b_i^\bullet$, then $\varphi_{\rho' \prec \rho}(b_\bullet^\bullet) = b_i^\bullet$ as well.

every curve carried by a diagonal extension of a large subtrack of τ'' is also carried by a diagonal extension of a large subtrack of τ' , in such a way that it runs along the branches b'_1 and b'_2 . Therefore, in order to verify that $\text{PN}(\tau'') \subseteq \text{int PN}(\tau)$, it suffices³ to check whether $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_1})$ and $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_2})$ fill a large subtrack of τ , which is straightforward to do in polynomial time in ξ . In other words, the existence of an intermediate train track τ' in which τ'' is deeply nested allows us to bypass the issue of having to iterate over (potentially too many) diagonal extensions of large subtracks of τ'' .

Implementation and complexity. Finally, we address the matter of how to represent and manipulate train tracks on a surface from an algorithmic point of view. We think of train tracks as abstract graphs, with vertices given by the switches and edges given by the branches; for each switch, we keep track of which branches are incident to each direction, as well as the order in which they appear along a co-orientation. Additionally, we store some information about the complementary regions, such as their topological type and the sequence of branches that appears along their boundary.

In short, we represent train tracks on S in such a way that we can reconstruct them up to homeomorphism of S , although not necessarily up to isotopy. While it would be theoretically possible to encode the specific way in which a train track is embedded in S – using a technique similar to that of normal curves – this information would be completely irrelevant for our purposes. In fact, all we need is a representation that is detailed enough to encode largeness of a train track, and supports the SPLIT and TWIST moves described in Section 3.1.

We remark that, for any sensible representation of train tracks, all the operations we will need to perform – and in particular the SPLIT and TWIST moves – can be implemented in polynomial time in the complexity ξ of the surface (recall that the number of switches and branches of a train track on S is $\mathcal{O}(\xi)$). However, the exact *degree* of this polynomial will depend on the specific data structure used to represent train tracks. Therefore, in order to avoid having to deal with implementation details,

³The reader will have surely realised that this is not an equivalent condition. To be more precise, the argument in the proof of Theorem 3.7 proceeds as follows. The intermediate train track τ' is chosen so that it is uniformly close to τ'' in $\mathcal{C}(S)$. Now, if $\varphi_{\tau' \prec \tau}(\nu)$ fills a large subtrack of τ – where ν is the fundamental measure on τ' given by Lemma 2.5 – then $\text{PN}(\tau'') \subseteq \text{int PN}(\tau)$ as described. Otherwise, it is easy to see that τ' is close to τ in $\mathcal{C}(S)$, and hence so is τ'' .

we will be content with stating a generic polynomial dependence on ξ of the running time of our algorithms. Instead, we will try and be more precise about the degree of the dependence on the complexities of curves and surface homeomorphisms.

A pre-measure μ on a train track τ is represented by a function $\mathcal{B}(\tau) \rightarrow \mathbb{Z}_{\geq 0}$; we define the *complexity* of μ to be

$$\|\mu\| = \sum_{b \in \mathcal{B}(\tau)} \log(\mu(b) + 1).$$

To describe a carrying $\tau' \prec \tau$ combinatorially, we simply keep track of the action of the carrying map $\varphi_{\tau' \prec \tau}$ on switches, orientations, and co-orientations of τ' , together with the pre-measures $\{\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'}) : b' \in \mathcal{B}(\tau')\}$ on τ . This information might not be enough to recover τ' from τ up to isotopy, but it is sufficient for computing the image $\varphi_{\tau' \prec \tau}(\mu') \in \mathcal{P}(\tau)$ of any measure $\mu' \in \mathcal{P}(\tau')$. The *complexity* of the carrying map $\varphi_{\tau' \prec \tau}$ is defined as

$$\|\varphi_{\tau' \prec \tau}\| = \sum_{b' \in \mathcal{B}(\tau')} \|\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'})\|.$$

3 Coarse distance algorithm

For a train track τ on a surface S , the set $\mathcal{F}(\tau)$ of fundamental curves of τ has universally bounded diameter in $\mathcal{C}(S)$, as it can be easily seen from Proposition 2.2. Therefore, we can think of train tracks on S as “coarse points” in $\mathcal{C}(S)$, and interpret sequences of train tracks as “coarse paths” in $\mathcal{C}(S)$.

Masur and Minsky show in [26, Theorem 1.3] that a splitting sequence $\tau_0 \succ \dots \succ \tau_n$ of train tracks, if considered as a coarse path in $\mathcal{C}(S)$, is an unparametrised quasi-geodesic, and in fact lies uniformly close to an actual geodesic. Bell and Webb [5] use this fact to construct a finite neighbourhood of this quasi-geodesic that is guaranteed to contain a (tight) geodesic between a fundamental curve of τ_0 and a fundamental curve of τ_n . Since the size of this neighbourhood is polynomial in n , they employ this construction to give a polynomial-time algorithm to compute distances in the curve graph.

The algorithm of Bell and Webb, however, does not run in polynomial time in the complexity of S . In this section, we present an algorithm that offers a trade-off between accuracy and good dependence of the running time on ξ . More precisely, given two essential curves a and b on S , our algorithm – as stated in Corollary 3.12 – computes an estimate d of the distance $\text{dist}_{\mathcal{C}(S)}(a, b)$ such that

$$d - L_+ \leq \text{dist}_{\mathcal{C}(S)}(a, b) \leq L_\times \cdot d + L_+.$$

The running time of the algorithm is polynomial in the parameters $\|a\|$, $\|b\|$, and ξ , and the constants L_\times and L_+ are polynomial in ξ .

The core idea of our algorithm is to use the combinatorial criterion of Proposition 2.3 to detect when the unparametrised quasi-geodesic given by $\tau_0 \succ \dots \succ \tau_n$ has made definite progress in $\mathcal{C}(S)$, and extract a subsequence corresponding to an actual quasi-geodesic parametrisation. The issue with this strategy is that, as already noted, the condition $\text{PN}(\tau') \subseteq \text{int PN}(\tau)$ is hard to verify in polynomial time in ξ for an arbitrary pair $\tau' \prec \tau$ of train tracks. To overcome this problem, we devise a procedure to construct a sequence of nested train tracks for which the condition above can be efficiently checked; as anticipated, the key property is that consecutive elements of the sequence should be deeply nested.

3.1 Agol–Hass–Thurston moves

Agol, Hass, and Thurston [1] introduce an algorithm to compute orbits of interval isometries in polynomial time; the algorithm works in complete generality, but the main application in their article is to efficiently compute properties of normal surfaces in 3-manifolds. We describe a specialised version of the Agol-Hass-Thurston algorithm for measured train tracks. This algorithm is simpler than the general one, in the sense that it involves only two types of moves, which can easily be interpreted in a purely topological fashion. The increased simplicity provides better running time bounds, and allows us to reason about combinatorial properties of the resulting train tracks.

We start by describing the two aforementioned moves. The input for both of them is a connectedly fully measured train track (τ, μ) – where we assume that τ is not a simple closed curve – and a switch s of τ , together with a co-orientation ω of s .

◇ The SPLIT move. Choose a direction η of s such that

$$\mu(\mathcal{E}_\eta(\tau)_1^\omega) \leq \mu(\mathcal{E}_{\bar{\eta}}(\tau)_1^\omega).$$

For ease of notation, let $e_i = \mathcal{E}_\eta(\tau)_i^\omega$. Moreover, denote by b_i the branch of which e_i is an end, and by \bar{b} the branch of which $\mathcal{E}_{\bar{\eta}}(\tau)_1^\omega$ is an end; we call this end \bar{b}^\bullet . Define

$$h = \max\{1 \leq i \leq |\mathcal{E}_\eta(\tau)| : \mu(b_1) + \dots + \mu(b_i) \leq \mu(\bar{b})\}.$$

We call the branches b_1, \dots, b_h *splitting branches*, and the ends e_1, \dots, e_h *splitting ends*.

The output of the SPLIT move is a measured train track $(\tau', \mu') \prec (\tau, \mu)$, obtained by pushing the ends e_1, \dots, e_h along \bar{b} all the way to the other end \bar{b}^\bullet . Figure 3.1 illustrates the construction of τ' , and shows how

$$\tau' = \tau \setminus (b_1 \cup \dots \cup b_h) \cup (b'_1 \cup \dots \cup b'_h),$$

where b'_i is a branch of τ' that follows $b_i \cup \bar{b}$ very closely. The carrying $\tau' \prec \tau$ satisfies the following properties: every non-splitting branch of τ is untouched in τ' , and for every splitting branch b_i of τ there is a unique branch b'_i of τ' such that $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_i}) \geq \mathbb{1}_{b_i}$. In fact, it is easy to see that $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_i})$ is either $\mathbb{1}_{b_i} + \mathbb{1}_{\bar{b}}$ or $\mathbb{1}_{b_i} + 2 \cdot \mathbb{1}_{\bar{b}}$, depending on whether b_i has one or two splitting ends.




Figure 3.1. The effect of applying a SPLIT move to a measured train track (τ, μ) with switch $s \in \mathcal{S}(\tau)$ and co-orientation ω of s . The measured train track (τ, μ) is represented by a thickened neighbourhood of τ , where the thickness of the branches is proportional to the weight assigned by μ . On the right, we only draw the train track $\tau' \prec \tau$, ignoring the measure μ' .

The measure μ' is chosen so that $\varphi_{\tau' \prec \tau}(\mu') = \mu$, and is obtained from μ as follows:

- if b is a non-splitting branch of τ different from \bar{b} , then $\mu'(b') = \mu(b)$, where b' is the branch of τ' such that $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'}) = \mathbb{1}_b$;
- $\mu'(\bar{b}') = \mu(\bar{b}) - \mu(b_1) - \dots - \mu(b_h)$, where \bar{b}' is the branch of τ' such that $\varphi_{\tau' \prec \tau}(\mathbb{1}_{\bar{b}'}) = \mathbb{1}_{\bar{b}}$;
- $\mu'(b'_i) = \mu(b_i)$ for every $1 \leq i \leq h$.

There is one exceptional case in which the construction is slightly different. If $\mu(\bar{b}) = \mu(b_1) + \dots + \mu(b_h)$, then the measure μ' would assign weight 0 to the branch \bar{b}' of τ' . To ensure that the measured train track (τ', μ') is connectedly fully measured, we remove the branch \bar{b}' from τ' , with the understanding that in this case the branch \bar{b} of τ is not untouched in τ' . In this case, we say that the branch \bar{b} *disappears* in τ' ; the situation is depicted in Figure 3.2. Moreover, when $\mu(\bar{b}) = \mu(s)$ – that is, when \bar{b} is the only branch incident to s in the direction $\bar{\eta}$ – not only does \bar{b} disappear in τ' , but there is no switch of τ' mapped to s by $\varphi_{\tau' \prec \tau}$. If this happens, we say that the switch s *disappears* in τ' .

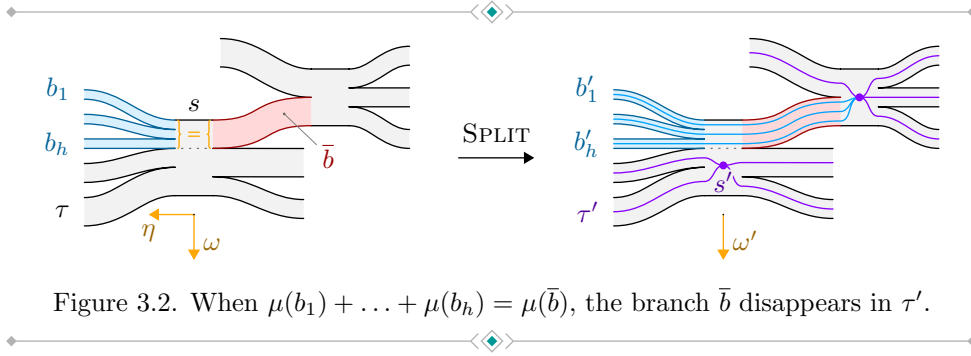


Figure 3.2. When $\mu(b_1) + \dots + \mu(b_h) = \mu(\bar{b})$, the branch \bar{b} disappears in τ' .

Without delving into implementation details, we remark that, from the description we have provided, it is easy to algorithmically construct the connectedly fully measured train track (μ', τ') and the carrying map $\varphi_{\tau' \prec \tau}$. We observe that the complexity of the carrying map is $\|\varphi_{\tau' \prec \tau}\| = \mathcal{O}(\xi)$, and the measure μ' satisfies $\|\mu'\| \leq \|\mu\|$. As long as a reasonable data structure is used to implement train tracks, the construction can be performed with running time $\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\|)$.

• The TWIST move. This move is an optimisation to perform multiple SPLIT moves in a row. In the setting above, suppose that $h < |\mathcal{E}_\eta(\tau)|$ and $\mathcal{E}_\eta(\tau)_{h+1}^\omega = \bar{b}_\bullet$. We still call the branches b_1, \dots, b_h *splitting branches*, and the ends e_1, \dots, e_h *splitting ends*. The output of the TWIST move is the measured train track $(\tau', \mu') \prec (\tau, \mu)$ obtained from k successive applications of the SPLIT move, where

$$k = \left\lfloor \frac{\mu(\bar{b})}{\mu(b_1) + \dots + \mu(b_h)} \right\rfloor.$$

To perform this move efficiently, we first note that the train track τ' is the result of applying $T_{\bar{b}}^{\pm k}$ to τ – which justifies the name “TWIST”; the construction is depicted in Figure 3.3. Just like in the SPLIT case, non-splitting branches of τ are untouched in τ' , and for every splitting branch b_i of τ there is a unique branch b'_i of τ' such that $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_i}) \geq \mathbb{1}_{b_i}$. More precisely, the pre-measure $\varphi_{\tau' \prec \tau}(\mathbb{1}_{b'_i})$ on τ is either $\mathbb{1}_{b_i} + k \cdot \mathbb{1}_{\bar{b}}$ or $\mathbb{1}_{b_i} + 2k \cdot \mathbb{1}_{\bar{b}}$, depending on whether b_i has one or two splitting ends.

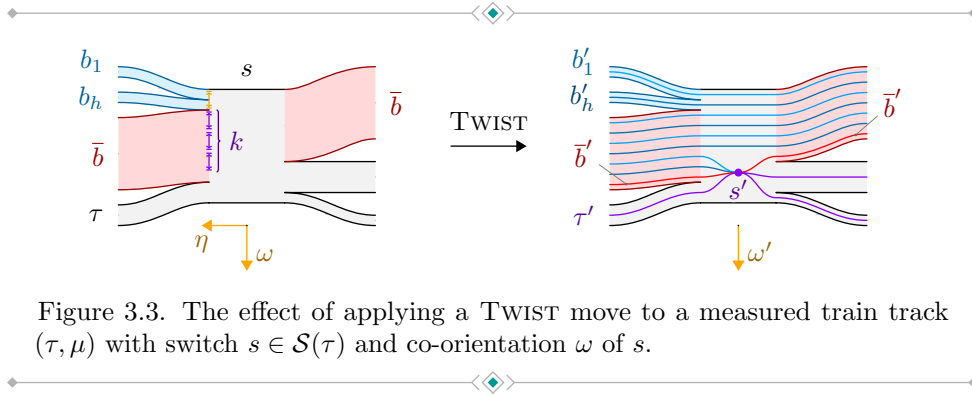


Figure 3.3. The effect of applying a TWIST move to a measured train track (τ, μ) with switch $s \in \mathcal{S}(\tau)$ and co-orientation ω of s .

The measure μ' on τ' is defined exactly like in the SPLIT case, the only difference being the formula

$$\mu'(\bar{b}') = \mu(\bar{b}) - k \cdot (\mu(b_1) + \dots + \mu(b_h)). \quad (2)$$

Just like in the SPLIT case, if the measure μ' would assign weight 0 to the branch \bar{b}' , we remove said branch from τ' , and say that \bar{b} *disappears* in τ' . In this case, in particular, the branch \bar{b} will not be untouched in τ' . Note that the switch s never disappears in τ' ; in other words, there is a unique switch s' of τ' such that $\varphi_{\tau' \prec \tau}(s') = s$.

The discussion above gives a complete recipe to construct the measured train track (τ', μ') and the carrying map $\varphi_{\tau' \prec \tau}$.⁴ We observe that the complexity of the carrying map is $\|\varphi_{\tau' \prec \tau}\| = \mathcal{O}(\xi \cdot \log k) = \mathcal{O}(\xi \cdot \|\mu\|)$, and the measure μ' satisfies $\|\mu'\| \leq \|\mu\|$. Moreover, the TWIST move can be performed with running time

$$\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\| \cdot \log \|\mu\|),$$

using the algorithm of Harvey and Hoeven [17] to compute the integers k and $\mu'(\bar{b}')$ in time $\mathcal{O}(\|\mu\| \cdot \log \|\mu\|)$.

Proximity. The first elementary property of the SPLIT and TWIST moves we observe is that they do not move train tracks too far in $\mathcal{C}(S)$. More precisely, we have the following bound.

Proposition 3.1 (Proximity). *If (τ', μ') is obtained from (τ, μ) by a SPLIT move or a TWIST move, then*

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq D_{\text{prox}},$$

where

$$D_{\text{prox}} = 17.$$

Proof. Let us employ the same notation we have used above when defining the two moves.

Suppose first that (τ', μ') is obtained from (τ, μ) by a SPLIT move. Let ν and ν' be fundamental measures on τ and τ' respectively. It is easy to see that $\varphi_{\tau' \prec \tau}(\nu')(z) \leq 2$ for every switch $z \in \mathcal{S}(\tau) \setminus \{s\}$, while

$$\varphi_{\tau' \prec \tau}(\nu')(s) = \nu'(s') + \nu'(b'_1) + \dots + \nu'(b'_h) \leq 4,$$

⁴In fact, since τ and τ' differ by a homeomorphism of S , they are represented by the same abstract graph.

where s' is the switch of τ' such that $\varphi_{\tau' \prec \tau}(s') = s$ – or, say, $\nu'(s') = 0$ if s disappears in τ' . As a consequence, we have the inequalities

$$i(\nu, \nu') \leq \sum_{z \in \mathcal{S}(\tau)} \nu(z) \cdot \varphi_{\tau' \prec \tau}(\nu')(z) \leq 4 |\mathcal{S}(\tau)| + 4 \leq 16\xi + 4,$$

which by Proposition 2.2 (and some easy algebra) imply that $\text{dist}_{\mathcal{C}(S)}(\nu, \nu') < 17$.

☞ Suppose now that (τ', μ') is obtained from (τ, μ) by a TWIST move. Note that the branches \bar{b} and \bar{b}' are isotopic simple closed curves. Moreover, every fundamental curve of τ (respectively τ') intersects \bar{b} (respectively \bar{b}') at most twice. This immediately gives the upper bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \bar{b}) + \text{diam}_{\mathcal{C}(S)}(\bar{b}', \mathcal{F}(\tau')) \leq 2 + 2 = 4. \quad \square$$

☞ **AHT sequences.** By iteratively applying SPLIT and TWIST moves, we obtain a sequence of nested measured train tracks. More precisely, let (τ, μ) be a connectedly fully measured train track, s a switch of τ , and ω a co-orientation of s . The *Agol–Hass–Thurston sequence* (or *AHT sequence* for short) generated by this data is a sequence of measured train tracks obtained from the following procedure.

We start by setting $(\tau_0, \mu_0) = (\tau, \mu)$, $s_0 = s$, and $\omega_0 = \omega$. On the i -th step, for $i \geq 0$, consider the measured train track (τ_i, μ_i) . If τ_i is a simple closed curve, then the procedure terminates. If we can apply a TWIST move to the switch s_i with co-orientation ω_i , then we do so, denoting by (τ_{i+1}, μ_{i+1}) the resulting measured train track. Otherwise, we apply a SPLIT move to the same co-oriented switch, and we call the resulting measured train track (τ_{i+1}, μ_{i+1}) . If the switch s_i disappears in τ_{i+1} , then we set $s_{i+1} = \emptyset$ and $\omega_{i+1} = \emptyset$, and the procedure terminates. Otherwise, we let s_{i+1} be the unique switch of τ_{i+1} such that $\varphi_{\tau_{i+1} \prec \tau_i}(s_{i+1}) = s_i$, and ω_{i+1} be the orientation of s_{i+1} such that $\varphi_{\tau_{i+1} \prec \tau_i}(\omega_{i+1}) = \omega_i$. We then proceed to the $(i+1)$ -th step.

This procedure will terminate after a finite number of steps, as Proposition 3.2 below implies. The output is a sequence of nested connectedly fully measured train tracks

$$(\tau_0, \mu_0) \succ (\tau_1, \mu_1) \succ \dots \succ (\tau_n, \mu_n),$$

where $(\tau_0, \mu_0) = (\tau, \mu)$, and either τ_n is the curve represented by μ , or τ_n has one less switch than τ . Note that, since the measured train tracks are generated one by one, we can stop the procedure prematurely whenever a condition of our choice is satisfied. The first m elements of an AHT sequence can be computed with running time

$$\mathcal{O}(\text{poly}(\xi) \cdot m \cdot \|\mu\| \cdot \log\|\mu\|).$$

There are a number of interesting objects associated to each measured train track (τ_i, μ_i) in an AHT sequence. We have already defined the switch s_i and the co-orientation ω_i in the description above; we now introduce the following additional notation for $0 \leq i < n$:

- we write $\varphi_{j,i}$ for $\varphi_{\tau_j \prec \tau_i}$ for every $i \leq j \leq n$;
- η_i is the direction of s_i such that the end $\mathcal{E}_{\eta_i}(\tau_i)_1^{\omega_i}$ is splitting in τ_i ; in particular, it satisfies

$$\mu_i(\mathcal{E}_{\eta_i}(\tau_i)_1^{\omega_i}) \leq \mu_i(\mathcal{E}_{\bar{\eta}_i}(\tau_i)_1^{\omega_i});$$

- \bar{b}_i is the branch of τ_i of which $\mathcal{E}_{\bar{\eta}_i}(\tau_i)_1^{\omega_i}$ is an end;
- if the branch \bar{b}_i disappears in τ_{i+1} , then we set $\bar{b}'_i = \emptyset$; otherwise, we let \bar{b}'_i be the unique branch of τ_{i+1} such that $\varphi_{i+1,i}(\mathbb{1}_{\bar{b}'_i}) = \mathbb{1}_{\bar{b}_i}$.

3.2 Elementary properties of AHT sequences

Let

$$(\tau_0, \mu_0) \succ (\tau_1, \mu_1) \succ \dots \succ (\tau_n, \mu_n)$$

be the AHT sequence generated by a connectedly fully measured train track (τ, μ) , a switch $s \in \mathcal{S}(\tau)$, and a co-orientation ω of s .

Carrying properties. Let us fix an integer $0 \leq i \leq n$. Figure 3.4 shows what the typical local picture of the carrying $\tau_i \prec \tau_0$ looks like near s .




Figure 3.4. A typical carrying $\tau_i \prec \tau_0$ between train tracks in an AHT sequence.

We now list a few properties of this carrying of train tracks that follow immediately from the procedure for generating the AHT sequence.

- 1 If a branch $b_0 \in \mathcal{B}(\tau_0)$ is untouched in τ_i , then there is a sequence of branches $b_1 \in \mathcal{B}(\tau_1)$, \dots , $b_i \in \mathcal{B}(\tau_i)$ such that $\varphi_{j,j-1}(\mathbb{1}_{b_j}) = \mathbb{1}_{b_{j-1}}$ for every $1 \leq j \leq i$. Moreover, the branches b_1, \dots, b_i are unique, and b_0, \dots, b_{i-1} are not splitting in $\tau_0, \dots, \tau_{i-1}$ respectively.

- 2 If B is any subset of $\mathcal{B}(\tau_i)$, then

$$|\{b \in \mathcal{B}(\tau_0) : \varphi_{i,0}(\mathbb{1}_{b'}) \geq \mathbb{1}_b \text{ for some } b' \in B\}| \geq |B|.$$

- 3 Suppose that τ_0 is one-switch. If e' is an end of τ_i and $e = \varphi_{i,0}(e')$, then

$$\text{height}_{\omega_0}(e) \leq \text{height}_{\omega_i}(e') \quad \text{and} \quad \text{height}_{\omega_i}(e') + \mu_i(e') \leq \text{height}_{\omega_0}(e) + \mu_0(e).$$

- 4 Suppose that τ_0 is one-switch. An end e of τ_0 is persistent in τ_i if and only if

$$\text{height}_{\omega_0}(e) < \mu_i(s_i).$$

In particular, for every direction η of s_0 , the set of ends in $\mathcal{E}_\eta(\tau_0)$ that are persistent in τ_i is a final segment of $\mathcal{E}_\eta(\tau_0)^{\omega_0}$.

- 5 Suppose that τ_0 is one-switch. Let b be a branch of τ_0 that is untouched in τ_i . Suppose that the ends b^\bullet and b_\bullet of b satisfy

$$\text{height}_{\omega_0}(b^\bullet) > \text{height}_{\omega_0}(b_\bullet).$$

Let b' be the (unique by Carrying property 1) branch of τ_i such that $\varphi_{i,0}(\mathbb{1}_{b'}) = \mathbb{1}_b$, and b'_\bullet be the end of b' such that $\varphi_{i,0}(b'_\bullet) = b_\bullet$. Then

$$\text{height}_{\omega_i}(b'^\bullet) > \text{height}_{\omega_i}(b'_\bullet),$$

and b'^\bullet is the unique end of τ_i whose image under $\varphi_{i,0}$ is b^\bullet .

Carrying property [4](#) can be proved as follows. If e is persistent in τ_i , then by Carrying property [3](#) we have

$$\text{height}_{\omega_0}(e) \leq \text{height}_{\omega_i}(e') < \mu_i(s_i),$$

where e' is any end of τ_i such that $\varphi_{i,0}(e') = e$. If, conversely, the end e satisfies $\text{height}_{\omega_0}(e) < \mu_i(s_i)$, then there is a unique end e' of τ_i such that $\text{dir}(e) = \text{dir}(\varphi_{i,0}(e'))$ and

$$\text{height}_{\omega_i}(e') \leq \text{height}_{\omega_0}(e) < \text{height}_{\omega_i}(e') + \mu_i(e').$$

Combining these inequalities with Carrying property [3](#) yields

$$\text{height}_{\omega_0}(\varphi_{i,0}(e')) \leq \text{height}_{\omega_0}(e) < \text{height}_{\omega_0}(\varphi_{i,0}(e')) + \mu_0(\varphi_{i,0}(e')),$$

which necessarily implies that $\varphi_{i,0}(e') = e$.

The other Carrying properties are immediate when $i = 1$, and can be proved by induction on i in the general case.

[D](#) **Exponential decay.** We now address the issue of how long an AHT sequence can be. It is quite clear that it has to be finite, since the quantity

$$\sum_{b \in \mathcal{B}(\tau_i)} \mu_i(b)$$

is a positive integer and strictly decreasing in i . However, the following proposition gives a much better bound than this first naive estimate.

Proposition 3.2 (Exponential decay). *Either $n \leq N_{\text{dec}}$, or there is an integer $1 \leq k \leq N_{\text{dec}}$ such that $\|\mu_k\| \leq \|\mu_0\| - 1$, where*

$$N_{\text{dec}} = 6\xi + 1.$$

Proof. Let us assume that $n > N_{\text{dec}}$; in particular, $s_i \neq \emptyset$ for every $0 \leq i \leq N_{\text{dec}}$.

[D](#) Suppose that $\bar{b}'_i = \emptyset$ for some $0 \leq i < N_{\text{dec}}$. Then

$$\|\mu_{i+1}\| = \|\mu_i\| - \log(\mu_i(\bar{b}_i) + 1) \leq \|\mu_i\| - 1,$$

and we can take $k = i + 1$. Consequently, from now on, we will assume that $\bar{b}'_i \neq \emptyset$ for every $0 \leq i < N_{\text{dec}}$.

[D](#) Note that, if $2\mu_{i+1}(\bar{b}'_i) < \mu_i(\bar{b}_i)$ for some $0 \leq i < N_{\text{dec}}$, then $\|\mu_{i+1}\| \leq \|\mu_i\| - 1$, and we can take $k = i + 1$. Therefore, from now on, we will assume that $2\mu_{i+1}(\bar{b}'_i) \geq \mu_i(\bar{b}_i)$ for every $0 \leq i < N_{\text{dec}}$.

[D](#) Suppose now that (τ_{i+1}, μ_{i+1}) is obtained from (τ_i, μ_i) by a TWIST move for some $0 \leq i < N_{\text{dec}}$. Then [\(2\)](#) implies that $\mu_{i+1}(\bar{b}'_i)$ is the remainder of $\mu_i(\bar{b}_i)$ modulo a positive integer smaller than $\mu_i(\bar{b}_i)$ – namely, the sum of the weights of the splitting ends of τ_i . In particular, we have that $2\mu_{i+1}(\bar{b}'_i) < \mu_i(\bar{b}_i)$, which contradicts our previous assumption. Therefore, from now on, we will assume that (τ_{i+1}, μ_{i+1}) is obtained from (τ_i, μ_i) by a SPLIT move for every $0 \leq i < N_{\text{dec}}$.

[D](#) Let us now fix an integer $0 \leq i < N_{\text{dec}} - 1$. Observe that, by definition of SPLIT move, the branch \bar{b}'_i is splitting in τ_{i+1} . By considering the SPLIT move from τ_{i+1} to τ_{i+2} , we see that

$$\mu_{i+2}(\bar{b}'_{i+1}) \leq \mu_{i+1}(\bar{b}_{i+1}) - \mu_{i+1}(\bar{b}'_i).$$

Combining this inequality with $2\mu_{i+1}(\bar{b}'_i) \geq \mu_i(\bar{b}_i)$ and $2\mu_{i+2}(\bar{b}'_{i+1}) \geq \mu_{i+1}(\bar{b}_{i+1})$ yields

$$\mu_i(\bar{b}_i) \leq \mu_{i+1}(\bar{b}_{i+1}) \quad \text{for every } 0 \leq i < N_{\text{dec}} - 1. \quad (3)$$

The intuition for why (3) gives a contradiction is that, since τ_0 has at most $6\xi < N_{\text{dec}}$ branches, two of the first N_{dec} SPLIT moves will operate on the same branch. More precisely, we claim the following: there are two indices $0 \leq i < j < N_{\text{dec}}$ such that $\varphi_{j,i}(\mathbb{1}_{\bar{b}_j}) \geq \mathbb{1}_{\bar{b}_i}$ as pre-measures on τ_i . In fact, if this were not the case, an easy induction using Carrying property 2 would show that the cardinality of the set

$$\{b \in \mathcal{B}(\tau_i) : \varphi_{j,i}(\mathbb{1}_{\bar{b}_j}) \geq \mathbb{1}_b \text{ for some } i \leq j < N_{\text{dec}}\}$$

is at least $N_{\text{dec}} - i$ for $0 \leq i < N_{\text{dec}}$, which is of course impossible for $i = 0$.

Let then $i < j$ be two indices as above. Since $\varphi_{j,i}(\mu_j) = \mu_i$, we immediately get that $\mu_j(\bar{b}_j) \leq \mu_i(\bar{b}_i)$, which we can refine to a strict inequality via the following argument. If $\varphi_{j,i+1}(\mathbb{1}_{\bar{b}_j}) \geq \mathbb{1}_{\bar{b}'_i}$, then $\mu_j(\bar{b}_j) \leq \mu_{i+1}(\bar{b}'_i) < \mu_i(\bar{b}_i)$; otherwise, $\mu_i(\bar{b}_i) \geq \mu_{i+1}(\bar{b}'_i) + \mu_j(\bar{b}_j)$. Either way, the strict inequality $\mu_j(\bar{b}_j) < \mu_i(\bar{b}_i)$ contradicts (3). \square

As a corollary, we observe that

$$i < N_{\text{dec}} \cdot (\|\mu_0\| - \|\mu_i\| + 1) \quad \text{for every } 0 \leq i \leq n. \quad (4)$$

3.3 Deep nesting algorithm

Our aim is now to describe an algorithm that, given as input a measured train track (τ, μ) , produces a measured train track $(\tau', \mu') \prec (\tau, \mu)$ such that τ' is deeply nested in τ and uniformly close to it in $\mathcal{C}(S)$. This algorithm will play a crucial role in our proof of Theorem 3.7. We start with a preliminary lemma.

Lemma 3.3. *Let (τ, μ) be a one-switch connectedly fully measured train track that is not a simple closed curve, and let*

$$(\tau_0, \mu_0) \succ \dots \succ (\tau_n, \mu_n)$$

be the AHT sequence generated by (τ, μ) , the unique switch s of τ , and a co-orientation ω of s . Let b be a branch of τ , and let b^\bullet be the end of b such that

$$\text{height}_\omega(b^\bullet) > \text{height}_\omega(b_\bullet).$$

Suppose that b is untouched in τ_{i-1} but not in τ_i for some $1 \leq i \leq n$. Then the end b^\bullet is not persistent in τ_i .

Proof. Let b' be the branch of τ_{i-1} such that $\varphi_{i-1,0}(\mathbb{1}_{b'}) = \mathbb{1}_b$ (which is unique by Carrying property 1), and let b'_\bullet be the end of b' such that $\varphi_{i-1,0}(b'_\bullet) = b_\bullet$. Carrying property 5 implies that

$$\text{height}_{\omega_{i-1}}(b'^\bullet) > \text{height}_{\omega_{i-1}}(b'_\bullet).$$

There are two possible reasons for b' to not be untouched in τ_i : either $b' = \bar{b}_{i-1}$ and it disappears in τ_i , or the end b'^\bullet is splitting in τ_{i-1} . In both cases, one checks that the end b'^\bullet is not persistent in τ_i ; since b'^\bullet is the only end of τ_{i-1} that $\varphi_{i-1,0}$ sends to b^\bullet , this implies that b^\bullet is not persistent in τ_i . \square

We now present an intermediate algorithm that, loosely speaking, takes the input train track “one step closer” to being deeply nested.

Proposition 3.4. *Let (τ, μ) be a one-switch connectedly fully measured train track, whose unique switch we denote by s . Let e and e' be two consecutive ends of τ . There is an algorithm to compute a one-switch connectedly fully measured train track $(\tau', \mu') \prec (\tau, \mu)$, a co-orientation ω of s , and a choice $e'' \in \{e, e'\}$ such that*

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq D'_{\text{d.n.}},$$

and e'' and all the ends of τ which are ω -smaller than e'' are not persistent in τ' . The constant $D'_{\text{d.n.}}$ is given by

$$D'_{\text{d.n.}} = D_{\text{prox}} + 10.$$

The running time of the algorithm is

$$\mathcal{O}(\text{poly}(\xi) \cdot (\|\mu\| - \|\mu'\| + 1) \cdot \|\mu\| \cdot \log\|\mu\|).$$

Proof. Note that τ cannot be a simple closed curve. Let b be the branch of which e is an end, and let us write b^\bullet for e .

🔗 Let us first suppose that b does not bound a punctured monogon in S . In this case, the desired co-orientation ω of s is the one such that

$$\text{height}_\omega(b^\bullet) > \text{height}_\omega(b_\bullet).$$

Let

$$(\tau_0, \mu_0) \succ \dots \succ (\tau_n, \mu_n)$$

be the AHT sequence generated by the measured train track (τ, μ) , the switch s , and the co-orientation ω . We set $(\tau', \mu') = (\tau_i, \mu_i)$, where

$$i = \min\{1 \leq j \leq n : b \text{ is not untouched in } \tau_j\}.$$

This is well-defined, because b is not untouched in τ_n – in fact, the train track τ_n is a simple closed curve. By Lemma 3.3, the end e is not persistent in τ_i . We are now left to show that τ_i is close to τ_0 in $\mathcal{C}(S)$. To see this, we refer the reader to Figure 3.5, which describes how to homotope b to a smooth curve c that intersects every fundamental curve of τ_0 and of τ_{i-1} at most twice. This gives the bound

$$\begin{aligned} \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), \mathcal{F}(\tau_i)) &\leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), c) + \text{diam}_{\mathcal{C}(S)}(c, \mathcal{F}(\tau_{i-1})) + \\ &\quad \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i-1}), \mathcal{F}(\tau_i)) \\ &\leq 2 + 2 + D_{\text{prox}} \leq D'_{\text{d.n.}}, \end{aligned}$$

where the bound on $\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i-1}), \mathcal{F}(\tau_i))$ comes from Proposition 3.1.

Clearly, the same argument works when e' is an end of a branch of τ that does not bound a punctured monogon in S .

🔗 Let ω' be the co-orientation of s such that $e <_{\omega'} e'$, and let \bar{e} be an end of τ such that $\text{dir}(\bar{e}) \neq \text{dir}(e)$ and

$$\text{height}_{\omega'}(\bar{e}) \leq \text{height}_{\omega'}(e) \leq \text{height}_{\omega'}(\bar{e}) + \mu(\bar{e}). \quad (5)$$

We now assume that \bar{e} is an end of a branch $\bar{b} \in \mathcal{B}(\tau)$ which does not bound a punctured monogon in S .




Figure 3.5. Since the branch b of τ_0 is untouched in τ_{i-1} , there is a smooth curve c homotopic to b that intersects every fundamental curve of τ_0 and of τ_{i-1} at most twice. When b is returning, as in (a), we can take $c = b$. Otherwise, we take c to be a copy of b that is smoothed out near s_0 , as shown in (b).

By the previous case, we can find a co-orientation ω of s and an integer i such that, if (τ_i, μ_i) is the i -th element of the AHT sequence generated by (τ, μ) , s , and ω , then

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), \mathcal{F}(\tau_i)) \leq D'_{\text{d.n}}$$

and \bar{e} is not persistent in τ_i . If $\omega' = \omega$, then Carrying property 4 and (5) imply that

$$\mu_i(s_i) \leq \text{height}_{\omega}(\bar{e}) \leq \text{height}_{\omega}(e)$$

and hence e is not persistent in τ_i either. If $\omega' = \bar{\omega}$, we use Carrying property 4 and (5) again to deduce that

$$\begin{aligned} \mu_i(s_i) &\leq \text{height}_{\omega}(\bar{e}) \\ &= \mu_0(s_0) - \text{height}_{\omega'}(\bar{e}) - \mu_0(\bar{e}) \\ &\leq \mu_0(s_0) - \text{height}_{\omega'}(e) \\ &= \text{height}_{\omega}(e'), \end{aligned} \tag{6}$$

hence e' is not persistent in τ_i .

We can then set $(\tau', \mu') = (\tau_i, \mu_i)$, and e'' to be e or e' depending on whether $\omega' = \omega$ or not.

🔗 The only case we are left to consider is when e , e' , and \bar{e} are all ends of branches of τ which bound punctured monogons in S . Let us write \bar{b}^\bullet for \bar{e} . Up to swapping e and e' , we can assume that $\bar{b}^\bullet \succ_{\omega'} \bar{b}_\bullet$. Set $\omega = \bar{\omega}'$, and let

$$(\tau_0, \mu_0) \succ \dots \succ (\tau_n, \mu_n)$$

be the AHT sequence generated by (τ, μ) , s , and ω . We set $(\tau', \mu') = (\tau_i, \mu_i)$, where

$$i = \min\{1 \leq j \leq n : b \text{ or } \bar{b} \text{ is not untouched in } \tau_j\}.$$

The curve represented by the measure $\mathbb{1}_b + \mathbb{1}_{\bar{b}}$ is a fundamental curve of both τ_0 and τ_{i-1} , hence it intersects every fundamental curve of τ_0 and of τ_{i-1} at most 4 times. The usual estimates then give the bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), \mathcal{F}(\tau_i)) \leq 5 + 5 + D_{\text{prox}} = D'_{\text{d.n}}.$$

Finally, we argue that one of e and e' is not persistent in τ_i . If b is not untouched in τ_i , then Lemma 3.3 implies that the ω -smaller end of b – which is either e or e' , since b bounds a punctured monogon in S – is not persistent in τ_i . If, instead, the branch \bar{b} is not untouched in τ_i , then we use Lemma 3.3 again to deduce that \bar{e} is not persistent in τ_i . In this case, the same computation carried out in (6) shows that e' is not persistent in τ_i .

Two final remarks are in order. Firstly, we have shown how to find a connectedly fully measured train track $(\tau', \mu') \prec (\tau, \mu)$, a co-orientation ω of s , and a choice $e'' \in \{e, e'\}$ such that the end e'' is not persistent in τ' . Since (τ', μ') comes from an AHT sequence, however, Carrying property 4 implies that all the ends of τ which are ω -smaller than e'' are also not persistent in τ' . Secondly, as far as the running time of the algorithm is concerned, note that we only need to perform the first i steps of the procedure generating the AHT sequence, where $(\tau', \mu') = (\tau_i, \mu_i)$. Each step can be performed in time $\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\| \cdot \log\|\mu\|)$, hence (4) gives the desired bound for the running time. \square

The final deep nesting algorithm is essentially a corollary of Proposition 3.4.

Proposition 3.5 (Deep nesting algorithm). *Let (τ, μ) be a one-switch connectedly fully measured train track. There is an algorithm to compute a one-switch connectedly fully measured train track $(\tau', \mu') \prec (\tau, \mu)$ such that τ' is deeply nested in τ and*

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq D_{\text{d.n.}},$$

where

$$D_{\text{d.n.}} = 12\xi \cdot D'_{\text{d.n.}}$$

The running time of the algorithm is

$$\mathcal{O}(\text{poly}(\xi) \cdot (\|\mu\| - \|\mu'\| + 1) \cdot \|\mu\| \cdot \log\|\mu\|)$$

Proof. The algorithm works as follows. Start by setting $(\tau_0, \mu_0) = (\tau, \mu)$. On the i -th step of the algorithm, for $i \geq 0$, we compute the set

$$E_i = \{e \in \mathcal{E}(\tau_0) : e \text{ is persistent in } \tau_i\}.$$

If $|E_i| = 2$ then the algorithm terminates and returns $(\tau', \mu') = (\tau_i, \mu_i)$. Otherwise, we can find two consecutive ends e and e' of τ_i such that $\varphi_{\tau_i \prec \tau_0}(e) \neq \varphi_{\tau_i \prec \tau_0}(e')$. Proposition 3.4 applied to the measured train track (τ_i, μ_i) and the ends e and e' yields a one-switch connectedly fully measured train track $(\tau_{i+1}, \mu_{i+1}) \prec (\tau_i, \mu_i)$ and a co-orientation ω of τ_i such that – without loss of generality – $e <_{\omega} e'$, and e and all the ends of τ_i which are ω -smaller than e are not persistent in τ_{i+1} . By Carrying property 4, it follows that $\varphi_{\tau_i \prec \tau_0}(e)$ is not persistent in τ_0 . In particular, the cardinality of E_{i+1} is strictly smaller than that of E_i . We then proceed to the $(i+1)$ -th step of the algorithm.

Suppose that the algorithm terminates after k steps – or, in other words, that $|E_k| = 2$. Clearly, the integer k is bounded above by $|\mathcal{E}(\tau)| - 2 < 12\xi$. Since

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau_{i+1})) \leq D'_{\text{d.n.}} \quad \text{for every } 0 \leq i < k,$$

we have the desired estimate

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq k \cdot D'_{\text{d.n.}} < D_{\text{d.n.}}$$

The bound on the running time follows from that of Proposition 3.4. \square

3.4 Distances for carried curves

Our next goal is to algorithmically estimate the distance in $\mathcal{C}(S)$ between a train track and a curve carried by it. As an intermediate result, we show that if a curve is carried by a train track τ , then it is also carried by a one-switch train track τ' which is uniformly close to τ in $\mathcal{C}(S)$.

Proposition 3.6 (One-switch algorithm). *Let (τ, μ) be a connectedly fully measured train track. There is an algorithm to compute a one-switch connectedly fully measured train track $(\tau', \mu') \prec (\tau, \mu)$ such that*

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq D_{1,s},$$

where

$$D_{1,s} = 48\xi^2 \cdot D_{\text{prox}}.$$

The running time of the algorithm is

$$\mathcal{O}(\text{poly}(\xi) \cdot (\|\mu\| - \|\mu'\| + 1) \cdot \|\mu\| \cdot \log\|\mu\|).$$

Proof. The algorithm works as follows. Start by setting $(\tau_0, \mu_0) = (\tau, \mu)$. On the i -step of the algorithm, for $i \geq 0$, choose a switch $s \in \mathcal{S}(\tau_i)$ arbitrarily. If s is the unique switch of τ_i , then the algorithm terminates with $(\tau', \mu') = (\tau_i, \mu_i)$. Otherwise, we distinguish two cases.

☞ If s has a returning branch b , then we arbitrarily choose another switch s' of τ_i and a co-orientation ω of s' . We let (τ_{i+1}, μ_{i+1}) be the last measured train track in the AHT sequence generated by (τ_i, μ_i) , s' , and ω . We note that τ_{i+1} has one less switch than τ_i , and moreover the branch b is untouched in τ_{i+1} . This implies that b is a curve whose intersection number with every fundamental curve of τ_i and of τ_{i+1} is at most 2. As a consequence, Proposition 2.2 gives the bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau_{i+1})) \leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), b) + \text{diam}_{\mathcal{C}(S)}(b, \mathcal{F}(\tau_{i+1})) \leq 4. \quad (7)$$

☞ If s has no returning branches, then we choose a co-orientation ω of s , and we let (τ_{i+1}, μ_{i+1}) be the last measured train track in the AHT sequence generated by (τ_i, μ_i) , s , and ω . Since s has no returning branches, it is easy to see that this AHT sequence only involves SPLIT moves, and its length is bounded above by the number of ends of τ_i which are incident to s . Therefore, by Proposition 3.1 we have the bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau_{i+1})) \leq |\mathcal{E}(\tau_i)| \cdot D_{\text{prox}} \leq 12\xi \cdot D_{\text{prox}}. \quad (8)$$

We then proceed to the $(i+1)$ -th step of the algorithm.

Suppose that the algorithm terminates after k steps – or, in other words, that τ_k is one-switch. Clearly, the integer k is bounded above by $|\mathcal{S}(\tau)| - 1 < 4\xi$. From (7) and (8) we get the desired inequality

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau')) \leq k \cdot 12\xi \cdot D_{\text{prox}} < D_{1,s}.$$

The bound on the running time follows from (4). □

The tools developed in this section, combined with Corollary 2.4, find their purpose in the proof of the following theorem.

Theorem 3.7 (Distance algorithm for carried curves). *Let (τ, μ) be a connectedly fully measured train track representing a curve $a \in \mathcal{C}_0(S)$. There is an algorithm to compute an integer $d \geq 0$ such that*

$$d - L_+^0 \leq \text{dist}_{\mathcal{C}(S)}(\mathcal{F}(\tau), a) \leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), a) \leq L_\times \cdot d + L_+^0,$$

where

$$\begin{aligned} L_+^0 &= D_{1.s} + 2D + 8, \\ L_\times &= 2D + 12, \\ D &= D_{d.n} + D_{\text{prox}} \cdot N_{\text{dec}}. \end{aligned}$$

The running time of the algorithm is

$$\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\|^2 \cdot \log\|\mu\|).$$

Proof. We split the algorithm in three steps.

🔗 **Multi-switch to one-switch.** The first step is to apply Proposition 3.6 to obtain a one-switch connectedly fully measured train track $(\tau_0, \mu_0) \prec (\tau, \mu)$ such that

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), \mathcal{F}(\tau_0)) \leq D_{1.s}. \quad (9)$$

If τ_0 is not large, then $\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), a) < D_{1.s} + 2 \leq L_+^0$, hence the algorithm can return $d = 0$. Therefore, from now on, we assume that the train track τ_0 is large.

🔗 **A sequence of nested train tracks.** We now compute a sequence of measured train tracks

$$(\tau_0, \mu_0) \succ (\tau'_0, \mu'_0) \succ (\tau_1, \mu_1) \succ (\tau'_1, \mu'_1) \succ \dots \succ (\tau'_{n-1}, \mu'_{n-1}) \succ (\tau_n, \mu_n) \succ (\bar{\tau}, \bar{\mu}), \quad (10)$$

according to the following procedure.

On the i -th step of the procedure, for $i \geq 0$, we apply the algorithm in Proposition 3.5 to (τ_i, μ_i) to find a one-switch connectedly fully measured train track $(\tau'_i, \mu'_i) \prec (\tau_i, \mu_i)$ such that τ'_i is deeply nested in τ_i and

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau'_i)) \leq D_{d.n}.$$

If τ'_i is not large, we set $n = i$ and $(\bar{\tau}, \bar{\mu}) = (\tau'_i, \mu'_i)$, and we terminate the procedure. Otherwise, we consider an AHT sequence generated by (τ'_i, μ'_i) , and let $(\tau_{i+1}, \mu_{i+1}) \prec (\tau'_i, \mu'_i)$ be the first measured train track in the sequence such that either τ_{i+1} is a simple closed curve, or $\|\mu_{i+1}\| \leq \|\mu'_i\| - 1$. By Proposition 3.1 and Proposition 3.2, we have the bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau'_i), \mathcal{F}(\tau_{i+1})) \leq D_{\text{prox}} \cdot N_{\text{dec}}.$$

If τ_{i+1} is not large, we set $n = i$ and $(\bar{\tau}, \bar{\mu}) = (\tau_{i+1}, \mu_{i+1})$, and we terminate the procedure. Otherwise, we proceed to the $(i + 1)$ -th step of the procedure.

When the procedure terminates, we are left with a sequence (10) of nested measured train tracks, of which we remark a few properties.

- Since $\|\mu_{i+1}\| \leq \|\mu'_i\| - 1$ for every $0 \leq i < n$, the length n is bounded above by $\|\mu\|$; in particular, this shows that we can compute the sequence (10) in time

$$\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\|^2 \cdot \log\|\mu\|).$$

- For $0 \leq i < n$, we have the uniform bound

$$\begin{aligned} \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau_{i+1})) &\leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_i), \mathcal{F}(\tau'_i)) + \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau'_i), \mathcal{F}(\tau_{i+1})) \\ &\leq D_{\text{d.n}} + D_{\text{prox}} \cdot N_{\text{dec}} = D. \end{aligned}$$

- Similarly, we have that $\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_n), \mathcal{F}(\bar{\tau})) \leq D$.
- The train tracks τ_i and τ'_i are large for $0 \leq i < n$, and so is τ_n .
- The train track $\bar{\tau}$ is not large, hence $\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\bar{\tau}), a) \leq 2$. As a consequence, we see that

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_n), a) \leq D + 2. \quad (11)$$

• **A quasi-geodesic subsequence.** Finally, we show how to compute a subsequence of (10) whose length is a coarse estimate for the distance in $\mathcal{C}(S)$ between τ and a . Start by setting $i_0 = 0$. On the j -th step of the procedure, for $j \geq 0$, we consider the set

$$I_j = \left\{ i_j < k \leq n : \begin{array}{l} \text{for every fundamental measure } \nu \text{ on } \tau_k, \text{ the} \\ \text{measure } \varphi_{\tau_k \prec \tau_{i_j}}(\nu) \text{ fills a large subtrack of } \tau_{i_j} \end{array} \right\}.$$

If $i_j = n$, then we set $m = j$ and terminate the procedure. If $i_j < n$ but I_j is empty, we set $i_{j+1} = n$ and $m = j + 1$, and terminate the procedure. Otherwise, we set $i_{j+1} = \min I_j$, and proceed to the $(j + 1)$ -th step.

In terms of computational complexity, note that, for a fixed integer $0 \leq k \leq n$, there are only $\mathcal{O}(\xi^2)$ fundamental measures on τ_k , and for each $\nu \in \mathcal{F}(\tau_k)$ we can compute the sets

$$\{b \in \mathcal{B}(\tau_h) : \varphi_{\tau_k \prec \tau_h}(\nu) \geq \mathbb{1}_b\} \quad \text{for } 0 \leq h \leq k$$

in time $\mathcal{O}(\text{poly}(\xi) \cdot k) = \mathcal{O}(\text{poly}(\xi) \cdot \|\mu\|)$. This information is enough to construct the subsequence $\{i_j : 0 \leq j \leq m\}$ which, given the sequence (10), can be done in time

$$\mathcal{O}(\text{poly}(\xi) \cdot \|\mu\|^2).$$

Our goal is now to prove that $d = \lfloor m/2 \rfloor$ is a valid return value.

- For the upper bound, fix an integer $1 \leq j \leq m$, and note that there is a fundamental curve ν of $\tau_{i_{j-1}}$ which is carried by a non-large subtrack σ of $\tau_{i_{j-1}}$. As a consequence, there is an essential curve $c \in \mathcal{C}_0(S)$ contained in the complement of σ , and we can choose c so that it intersects every branch of $\tau_{i_{j-1}}$ at most twice. In particular, the curve c intersects every fundamental curve of $\tau_{i_{j-1}}$ at most 4 times, and Proposition 2.2 then gives the bound

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i_{j-1}}), \nu) \leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i_{j-1}}), c) + \text{dist}_{\mathcal{C}(S)}(c, \nu) \leq 5 + 1 = 6,$$

which in turn yields

$$\begin{aligned} \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i_{j-1}}), \mathcal{F}(\tau_{i_j})) &\leq \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_{i_{j-1}}), \nu) + \text{diam}_{\mathcal{C}(S)}(\nu, \mathcal{F}(\tau_{i_j})) \\ &\leq 6 + D. \end{aligned}$$

We conclude that

$$\text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), \mathcal{F}(\tau_n)) \leq (D + 6) \cdot m \quad (12)$$

◻ For the lower bound, we start by proving that $\text{PN}(\tau'_{i_j}) \subseteq \text{int PN}(\tau_{i_{j-1}})$ for every $1 \leq j < m$. Let σ' be a large recurrent subtrack of τ'_{i_j} , and let ρ' be a recurrent diagonal extension of σ' . The carrying $\tau'_{i_j} \prec \tau_{i_j}$ induces a carrying $\sigma' \prec \sigma$ for a large recurrent subtrack σ of τ_{i_j} ; by choosing the smallest possible such σ , we can assume that σ' fills σ . Since τ'_{i_j} is deeply nested in τ_{i_j} , it is clear that σ' is deeply nested in σ , the two persistent ends of σ in σ' being exactly the two persistent ends of τ_{i_j} in τ'_{i_j} . Lemma 2.5 then implies that ρ' is carried by and deeply nested in a recurrent diagonal extension ρ of σ , in such a way that the two persistent ends of ρ in ρ' are ends of branches of σ . Moreover, Lemma 2.5 applied to the carrying $\rho' \prec \rho$ shows that there is a fundamental measure ν on ρ such that $\varphi_{\rho' \prec \rho}(\mu') \geq \nu$ for every non-zero $\mu' \in \mathcal{P}(\rho')$, and in fact the measure ν is supported on σ . We refer the reader to Figure 3.6 for a graphical depiction of the carryings $\sigma' \prec \sigma$ and $\rho' \prec \rho$.




Figure 3.6. (a) The subtrack σ' is carried by and deeply nested in σ ; the two ends of σ which are persistent in σ' are highlighted. (b) The diagonal extension ρ' of σ' is carried by and deeply nested in a diagonal extension ρ of σ . The two ends of ρ which are persistent in ρ' are ends of branches of σ .

The train track σ fills a large subtrack σ_* of $\tau_{i_{j-1}}$, and ρ is carried by a large recurrent diagonal extension ρ_* of σ_* . If μ' is any non-zero measure on ρ' , the inequality $\varphi_{\rho' \prec \rho}(\mu') \geq \nu$ implies that $\varphi_{\rho' \prec \rho_*}(\mu') \geq \varphi_{\rho \prec \rho_*}(\nu)$. Since $\varphi_{\tau_{i_j} \prec \tau_{i_{j-1}}}(\nu)$ fills a large subtrack of $\tau_{i_{j-1}}$, we see that $\varphi_{\rho' \prec \rho_*}(\mu')$ fills a diagonal extension of a large subtrack of $\tau_{i_{j-1}}$. In other words, we have proved that $\text{PN}(\tau'_{i_j}) \subseteq \text{int PN}(\tau_{i_{j-1}})$, and in particular that $\text{PN}(\tau_{i_{j+1}}) \subseteq \text{int PN}(\tau_{i_{j-1}})$. By means of Corollary 2.4 and (1), we finally get the lower bound

$$\text{dist}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0), \mathcal{F}(\tau_n)) \geq \left\lfloor \frac{m}{2} \right\rfloor - \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau_0)) \geq \left\lfloor \frac{m}{2} \right\rfloor - 5. \quad (13)$$

We now combine (9), (11), (12) and (13) to get

$$\begin{aligned} \text{diam}_{\mathcal{C}(S)}(\mathcal{F}(\tau), a) &\leq (D+6) \cdot m + D_{1,s} + D + 2, \\ \text{dist}_{\mathcal{C}(S)}(\mathcal{F}(\tau), a) &\geq \left\lfloor \frac{m}{2} \right\rfloor - D_{1,s} - D - 7. \end{aligned}$$

Therefore, $d = \lfloor m/2 \rfloor$ is a valid return value. ◻

3.5 Distances in the curve graph

The main application of Theorem 3.7 is to compute coarse distances between curves in the curve graph. In particular, we will show in Proposition 3.9 how to estimate the

distance between a very “short” curve and any other curve. A theorem of Lackenby and Yazdi [23] will then allow us to generalise this result to coarsely compute distances between any two curves. We start by describing how to find a train track carrying an arbitrary normal curve on a surface S with a triangulation \mathcal{T} . This task is essentially trivial when S is punctured, but requires some effort in the closed case.

Proposition 3.8. *Let $a \in \mathcal{C}_0(S)$ be a curve on S . There is an algorithm to find a train track τ and a measure $\mu \in \text{int P}(\tau)$ representing a such that every edge of \mathcal{T} intersects τ transversely and at most once. The running time of the algorithm is polynomial in ξ and $\|a\|$.*

Proof. In the context of this proof, by *pre-track* we mean a smooth 1-complex embedded in S which satisfies all the properties of a train track except for those concerning complementary regions.

When S is punctured, the argument is standard. We let τ_0 be the pre-track obtained by repeating the pattern shown in Figure 3.7a for each triangle of \mathcal{T} . The pre-track τ_0 has $|\mathcal{T}|$ switches, all of which are 4-valent, and three branches per triangle of \mathcal{T} . We can define a measure μ_0 on τ_0 by counting the number of normal arcs that run parallel to each branch, as shown in Figure 3.7b. By removing branches of measure zero from τ_0 , we get a train track τ and an induced measure $\mu \in \text{int P}(\tau)$ representing a .

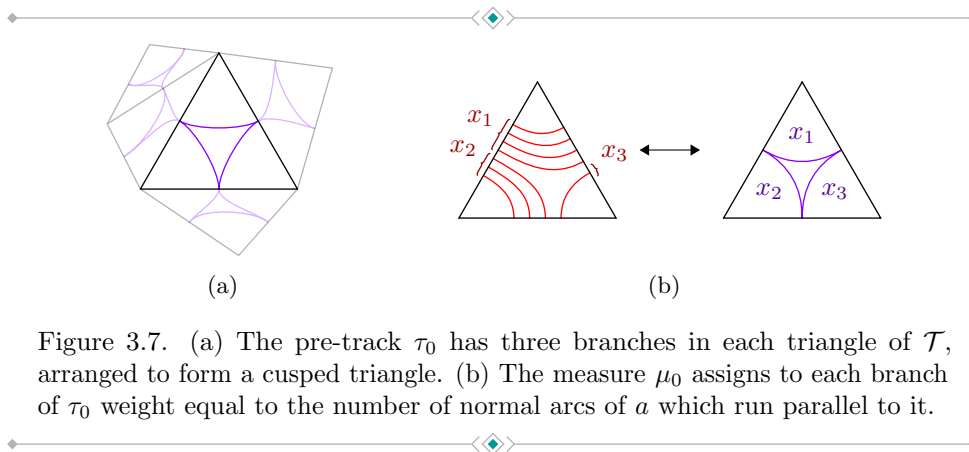


Figure 3.7. (a) The pre-track τ_0 has three branches in each triangle of \mathcal{T} , arranged to form a cusped triangle. (b) The measure μ_0 assigns to each branch of τ_0 weight equal to the number of normal arcs of a which run parallel to it.

In general, let us call the pre-track τ obtained from the procedure above the *pre-track associated* to the normal curve a . In the closed case, there is no guarantee that the pre-track associated to a is an actual train track, because it may have monogon or bigon complementary regions. Therefore, we need to preemptively modify the normal curve a , by possibly isotopying it across the vertex v of \mathcal{T} .

Let us number the types of normal arcs in \mathcal{T} by $1, 2, \dots, n$ in the order in which they appear counter-clockwise around v ; here $n = 2|\mathcal{T}|$ is three times the number of triangles of \mathcal{T} . We will call a normal arc of a *innermost* if it is the closest to v amongst all normal arcs of the same type, and *outermost* if it is the farthest away. For a normal curve b and an integer $1 \leq i \leq n$, let us call $x(b)_i$ the number of normal arcs of b of type i . Clearly, the integers $x(b)_1, \dots, x(b)_n$ are enough to recover the vector $\mathcal{T}(b)$, and vice versa.

We represent the triangles of \mathcal{T} as they appear around v , as shown in Figure 3.8. For $1 \leq i \leq n$, let Δ_i be the triangle of \mathcal{T} containing the normal arc type i . We remark that, in this representation, every triangle appears three times, but a triangle can never be

adjacent to itself; in other words, $\Delta_i \neq \Delta_{i+1}$ for $1 \leq i < n$, and $\Delta_n \neq \Delta_1$. This follows from the fact that, in a one-vertex triangulation of a surface, a triangle cannot be glued to itself along two of its edges.




Figure 3.8. We draw the triangles of \mathcal{T} in the order they appear counter-clockwise around the vertex v of \mathcal{T} . These pictures show what the pre-track associated to a looks like when it has a monogon (a) or bigon (b) complementary region.

We claim that the pre-track τ_0 associated to a normal curve b has a monogon (respectively bigon) complementary region if and only if $x(b)_i = 0$ for exactly one (respectively two) integers $1 \leq i \leq n$. The reverse implication is clear, as shown in Figure 3.8a (respectively Figure 3.8b). For the forward implication, suppose that a component Z of $S \setminus \tau_0$ is a monogon (respectively bigon). Since no complementary region entirely contained in a triangle of \mathcal{T} is a monogon (respectively bigon), the vertex v must lie in Z ; since no edge can be entirely contained in a monogon (respectively bigon), every edge of \mathcal{T} intersects τ_0 exactly once. It follows that, if $x(b)_i = 0$ for some integer $1 \leq i \leq n$, then $x(b)_j$ and $x(b)_k$ are both positive, where j and k refer to the two types of normal arcs other than i which appear in Δ_i . Therefore, the region Z is a cusped disc, with as many cusps as there are integers $1 \leq i \leq n$ such that $x(b)_i = 0$.

We now describe how to modify the normal curve a to ensure that τ_0 is an actual train track.

🔗 In order to get rid of monogons, we begin by applying [22, Theorem 6.3] to put a in minimal position with an arbitrary edge e of \mathcal{T} ; for ease of notation, we keep calling the resulting normal curve a . Since a is essential, we can assume that $x(a)_1 = 0$. Suppose, for the sake of contradiction, that the pre-track associated to a has a monogon complementary region; in other words, assume that $x(a)_i \geq 1$ for every $2 \leq i \leq n$. This situation is depicted in Figure 3.9a.

Let $2 \leq h, k \leq n$ be the integers such that the normal arc types 1, h , and k appear counter-clockwise in Δ_1 . We observe that, since $\Delta_2 \neq \Delta_1$, we have that $x(a)_k = x(a)_{k-1} + x(a)_2 \geq 2$. Similarly, the integer $x(a)_h$ is also at least 2. Let α be the subarc of a obtained by taking the union of the innermost normal arcs of types 2, \dots , n and the outermost normal arcs of types k and h . Since $x(a)_h \geq 2$ and $x(a)_k \geq 2$, we see that these normal arcs are all different, and hence the arc α is embedded in S . We can then isotope α across v , fixing $a \setminus \alpha$, to reduce the number of intersections with e , as depicted in Figure 3.9b. It follows that a was not in minimal position with e .

🔗 Suppose now that the pre-track associated to a has a bigon complementary region. This situation is depicted in Figure 3.10a. There is a unique integer $2 \leq j \leq n$ such




Figure 3.9. If the pre-track associated to a has a monogon complementary region (a), then a is not in minimal position with e , since it can be isotoped to reduce the number of intersections with it (b).

that $x(a)_j = 0$, and we can assume that $2 \leq j \leq n/2 + 1$. Let $1 \leq l, m \leq n$ be the integers such that the normal arc types j, l , and m appear counter-clockwise in Δ_j .

We remark that, if $j < h$, then we have that $\{h+1, n\} \cap \{1, j\} = \emptyset$, and hence $x(a)_h = x(a)_{h+1} + x(a)_n \geq 2$. Similarly, if $j < m$, then $x(a)_m = x(a)_{m-1} + x(a)_{j+1} \geq 2$.

Let α be the subarc of a obtained by taking the union of the innermost normal arcs of types $j+1, \dots, n$ and the outermost normal arcs of types h and m . By the remark above, all these normal arcs are different, and the arc α is embedded in S . We can then isotope α across v , fixing $a \setminus \alpha$, to obtain a new normal curve a' , as shown in Figure 3.10b. The numbers $x(a')_1, \dots, x(a')_n$ can easily be computed in terms of $x(a)_1, \dots, x(a)_n$ as follows. For each $1 \leq i \leq n$, define

$$r(i) = |\{i\} \cap \{2, \dots, j-1\}| + |\{i\} \cap \{k, l\}| - |\{i\} \cap \{j+1, \dots, n\}| - |\{i\} \cap \{h, m\}|.$$

Then $x(a')_i = x(a)_i + r(i)$. In particular, note that $x(a')_1 = x(a')_j = 0$. This means that either $x(a')_i = 0$ for at least three integers $1 \leq i \leq n$, or we can repeat this arc isotopy procedure with a' instead of a . By induction, if we let

$$t = \min \left(\{x(a)_i : 1 \leq i \leq n, r(i) = -1\} \cup \left\{ \left\lfloor \frac{x(a)_i}{2} \right\rfloor : 1 \leq i \leq n, r(i) = -2 \right\} \right),$$

then there is a normal curve a'' with $x(a'')_i = x(a)_i + t \cdot r(i)$, obtained by iterating the arc isotopy procedure t times. This curve satisfies $x(a'')_1 = x(a'')_j = 0$, and $x(a'')_i = 0$ for a third integer $2 \leq i \leq n$. In particular, the normal curve a'' is isotopic to a , and its associated pre-track has no monogon or bigon complementary regions.

Two concluding remarks in the closed case are in order. Firstly, note that the curve a'' – or, more precisely, the integers $x(a'')_1, \dots, x(a'')_n$ – can be computed in one single step, without actually performing the arc isotopy procedure $t = \mathcal{O}(2^{\|a\|})$ times. This guarantees that the running time of the algorithms remains polynomial in $\|a\|$. Secondly, for the bigon removal argument to work, it is crucial that the genus of S is at least 2. In fact, this condition guarantees that $n > 6$ and hence that $r(i) \in \{-1, -2\}$ for some $1 \leq i \leq n$, thus making t well-defined. When S is a torus, it may happen that $a' = a$ as normal curves, thus rendering the arc isotopy procedure ineffective. \square




Figure 3.10. If the pre-track associated to a has a bigon complementary region (a), we can isotope a subarc α of a across the vertex of the triangulation (b).

As anticipated above, Proposition 3.8 immediately gives an algorithm to estimate the distance in $\mathcal{C}(S)$ between two curves a and b , one of which (a in the statement below) is short.

Proposition 3.9 (Distance algorithm for short curves). *Let $a, b \in \mathcal{C}_0(S)$ be curves on S . There is an algorithm to compute an integer $d \geq 0$ such that*

$$d - \frac{6\|a\|_1 + 6}{\log \xi} - L_+^0 - 2 \leq \text{dist}_{\mathcal{C}(S)}(a, b) \leq L_\times \cdot d + \frac{6\|a\|_1 + 6}{\log \xi} + L_+^0 + 2.$$

The running time of the algorithm is polynomial in ξ and $\|b\|$.

Proof. We use Proposition 3.8 to compute a connectedly fully measured train track (τ, μ) representing b , such that τ intersect every edge of \mathcal{T} at most once. Note that, as a consequence, every fundamental curve of τ can be realised so that it intersects a at most

$$\sum_{e \in \mathcal{T}} 2\mathcal{T}(a)_e$$

times. Proposition 2.2 immediately implies that

$$\text{diam}_{\mathcal{C}(S)}(a, \mathcal{F}(\tau)) \leq \frac{6\|a\|_1 + 6}{\log \xi} + 2.$$

It is then clear that the integer d returned by the algorithm in Theorem 3.7 applied to (τ, μ) satisfies the desired inequalities. \square

Remark 3.10. The statement of Proposition 3.9 claims that the running time of the algorithm is polynomial in $\|b\|$, without any estimate on the degree of the polynomial. The reason for this omission is that the proof relies on the algorithm of Proposition 3.8, which in turn uses [22, Theorem 6.3] in the case where S is closed; Lackenby does not give any upper bound to the degree of the polynomial in the running time of the latter.

However, we can be more precise if we assume that S is punctured. In this case, the algorithm of Proposition 3.8 runs in linear time in $\|b\|$. As a consequence, the running time of the algorithm in Proposition 3.9 is the same as that of Theorem 3.7, that is

$$\mathcal{O}(\text{poly}(\xi) \cdot \|b\|^2 \cdot \log \|b\|). \quad \times$$

In order to get rid of the dependence on $\|a\|_1$ of the estimate in Proposition 3.9, we use an algorithm of Lackenby and Yazdi to find a triangulation of S with respect to which a is very short. We remark that a similar result was already proved by Bell [6], but without the polynomial dependence on the complexity of S .

Theorem 3.11 ([23, Theorems 7.4 and 7.5]). *Let $a \in \mathcal{C}_0(S)$ be a curve. There is an algorithm to produce a path in the flip and twist graph of S between \mathcal{T} and a triangulation \mathcal{T}' of S such that a intersects every edge of \mathcal{T}' at most twice. The length of the path and the running time of the algorithm are polynomial in ξ and $\|a\|_{\mathcal{T}}$.*

Note that Theorems 7.4 and 7.5 of [23] are not stated in terms of flip and twist graph. However, the sequences of triangulations of S that Lackenby and Yazdi construct can be interpreted as paths in $\mathcal{G}(S)$, where two consecutive triangulations are joined by an arc of length $\mathcal{O}(\log(\xi) + \|a\|_{\mathcal{T},1})$. Since the lengths of their sequences and the running times of their algorithms are polynomial in ξ and $\|a\|_{\mathcal{T}}$, this justifies the statement of Theorem 3.11. We also remark that, for the closed case, Theorem 7.4 of [23] produces a triangulation \mathcal{T}' of which a is an edge. However, every edge of \mathcal{T}' can be isotoped to a normal curve intersecting every (other) edge at most twice, as we require in Theorem 3.11.

By combining Proposition 3.9 and Theorem 3.11, we finally have a general algorithm to estimate distances in the curve graph.

Corollary 3.12 (Distance algorithm). *Let $a, b \in \mathcal{C}_0(S)$ be curves. There is an algorithm to compute an integer $d \geq 0$ such that*

$$d - L_+ \leq \text{dist}_{\mathcal{C}(S)}(a, b) \leq L_{\times} \cdot d + L_+,$$

where

$$L_+ = L_+^0 + 36.$$

The running time of the algorithm is polynomial in ξ , $\|a\|$, and $\|b\|$.

Proof. We use Theorem 3.11 to find a path in $\mathcal{G}(S)$ between \mathcal{T} and a triangulation \mathcal{T}' , every edge of which intersects a at most twice. The path has polynomial length in ξ and $\|a\|_{\mathcal{T}}$, and we use it to compute the vectors $\mathcal{T}'(a)$ and $\mathcal{T}'(b)$. We observe that the sum of the entries of $\mathcal{T}'(a)$ is at most $2|\mathcal{T}'| \leq -6\chi + 6$; as a consequence, some easy algebra shows that

$$\frac{6\|a\|_{\mathcal{T}',1} + 6}{\log \xi} \leq 34.$$

Moreover, note that the complexity $\|b\|_{\mathcal{T}'}$ is polynomial in $\|b\|_{\mathcal{T}}$ and in the length of the path in $\mathcal{G}(S)$ from \mathcal{T} to \mathcal{T}' ; in particular, we have that

$$\|b\|_{\mathcal{T}'} = \mathcal{O}(\text{poly}(\xi) \cdot \text{poly}(\|a\|_{\mathcal{T}}) \cdot \text{poly}(\|b\|_{\mathcal{T}})).$$

Finally, we apply Proposition 3.9 to the curves a and b with respect to the triangulation \mathcal{T}' ; the estimate d that we get satisfies the required bounds for this algorithm. \square

Remark 3.13. For the sake of simplicity, we have stated Theorem 3.7 as a coarse distance algorithm. However, it is clear from the proof that the same algorithm can be used to produce an explicit quasi-geodesic between a fundamental curve ν of τ and a . In fact, using the notation from the proof, we can arbitrarily choose a fundamental curve ν_j of τ_{i_j} for each $0 \leq j \leq m$. It is then easy to see that $\nu, \nu_0, \dots, \nu_m, a$ is a (L_{\times}, L_+^0) -quasi-geodesic from ν to a , in the sense that

$$|j - k| - L_+^0 \leq \text{dist}_{\mathcal{C}(S)}(\nu_j, \nu_k) \leq L_{\times} \cdot |j - k| + L_+^0 \quad \text{for } -1 \leq j, k \leq m + 1,$$

where $\nu_{-1} = \nu$ and $\nu_{m+1} = a$.

Note that, since all the carrying maps in the sequence (10) are computed explicitly, we can think of the curves $\nu_{-1}, \dots, \nu_{m+1}$ as being carried by τ . It is then straightforward to go through the proofs of Proposition 3.9 and Corollary 3.12 and obtain an algorithm to compute a (L_\times, L_+) -quasi-geodesic between two given curves $a, b \in \mathcal{C}_0(S)$. \times

4 Nielsen-Thurston classification

A remarkable theorem of Thurston (see [13]) states that every mapping class in $\text{Mcg}(S)$ is either periodic, reducible (and non-periodic), or pseudo-Anosov; the category to which a mapping class belongs is called its Nielsen-Thurston type. We will say that a homeomorphism of S is periodic, reducible, or pseudo-Anosov if its isotopy class in $\text{Mcg}(S)$ is. In this section, we endeavour to describe an efficient algorithm to decide the Nielsen-Thurston type of a surface homeomorphism, by studying its action on the curve graph of S and, in particular, its stable translation length (see Section 4.3). Naturally, the distance algorithm devised in Section 3 will play a crucial role here.

The argument to turn a distance algorithm for the curve graph into a surface homeomorphism classification algorithm is not novel. In fact, this section closely follows [5, §4.2]. The three main differences are that we write down the proofs in greater detail, we make an effort to keep the constants and the running time bounds polynomial in ξ , and we adapt the arguments to make use of a coarse distance algorithm rather than an exact one. We also remark that, unlike Bell and Webb's, our algorithm does not produce a fixed multicurve in the reducible case.

4.1 Periodic homeomorphisms

Of the three Nielsen-Thurston types, periodic homeomorphisms are the easiest to identify, thanks to the following bound on their order.

Proposition 4.1. *If a homeomorphism $f: S \rightarrow S$ is periodic, then f^k is isotopic to the identity for some $1 \leq k \leq C_{\text{per}}$, where*

$$C_{\text{per}} = -6\chi.$$

Proof. If S is closed and has genus g , then it was proved by Wiman [35] that the order of f in $\text{Mcg}(S)$ is bounded above by $4g+2$, which is smaller than -6χ for $g \geq 2$. Otherwise, let f be a homeomorphism of the punctured surface S with finite order k . By the Nielsen realisation theorem [27], we can assume that f is in fact an isometry of some hyperbolic metric on S , such that f^k is the identity. If we let H be the cyclic group of isometries of S generated by f , then the quotient $S' = S/H$ is a topological surface, and the projection $\pi: S \rightarrow S'$ is a branched covering. The Riemann-Hurwitz formula implies that

$$-\frac{\chi(S)}{k} = -\chi(S') + \sum_{x \in S'} \left(1 - \frac{1}{e_x}\right),$$

where e_x is the cardinality of the stabilisers in H of the preimages of x under π . Note that, since S is punctured, then so is S' , and hence $\chi(S') \leq 1$. A simple case analysis shows that, if the right-hand side of the equality is positive, then it is not smaller than $1/6$. The desired inequality $k \leq -6\chi(S)$ follows immediately. \square

Thanks to the bound of Proposition 4.1, we can use the Alexander method described in [12, §2.3] to detect periodic homeomorphisms.

Proposition 4.2. *Let $f: S \rightarrow S$ be a homeomorphism. There is an algorithm to decide if f is periodic. The running time of the algorithm is polynomial in ξ and $\|f\|$.*

Proof. We show how to construct a small set \mathcal{A} of short curves on S such that if a homeomorphism $S \rightarrow S$ fixes each curve in \mathcal{A} up to isotopy then it has finite order. By Proposition 4.1, this is enough to recognise periodic homeomorphisms.

☞ Suppose first that S is punctured, and recall that the triangulation \mathcal{T} is ideal. In this case, we take \mathcal{A} to be

$$\mathcal{A} = \{a : e \text{ is an edge of } \mathcal{T}, a \text{ is boundary component of a neighbourhood of } e \text{ in } S\},$$

where we think of punctures of S as marked points on a closed surface. The curves in \mathcal{A} are easy to construct algorithmically, and have complexity $\mathcal{O}(\xi)$; moreover, $|\mathcal{A}| = \mathcal{O}(\xi)$. Finally, we note that a homeomorphism $S \rightarrow S$ fixes all the curves in \mathcal{A} up to isotopy if and only if it fixes all the edges of \mathcal{T} up to isotopy. By the Alexander method [12, Proposition 2.8], a homeomorphism $S \rightarrow S$ fixing the isotopy classes of the edges of an ideal triangulation of S is isotopic to the identity.

☞ The closed case is more involved. Let e_1, \dots, e_m be the edges of the one-vertex triangulation \mathcal{T} . We start by considering the set of curves

$$\mathcal{A}_1 = \left\{ a : \begin{array}{l} 1 \leq i \leq m, a \text{ is a boundary component of a} \\ \text{neighbourhood of } e_1 \cup \dots \cup e_i \text{ in } S \text{ and is essential} \end{array} \right\}.$$

We can assume that the curves in \mathcal{A}_1 are pairwise non-isotopic, since Proposition 2.1 provides a procedure to remove duplicates.

The elements of \mathcal{A}_1 can be constructed by tracing the boundary of a neighbourhood of $e_1 \cup \dots \cup e_i$ and then normalising the resulting curves. Since these curves are short – they intersect the triangulation $\mathcal{O}(\xi)$ times – the normalising procedure can be performed directly, without resorting to the machinery of [22, Theorem 1.3].

By [23, Lemma 3.8], the curves in \mathcal{A}_1 form a pants decomposition of the surface S . There is a multicurve A_2 on S which intersects every pair of pants in three arcs as shown in Figure 4.1. This multicurve can be constructed algorithmically by tracing three arcs in each pair of pants, then joining arcs in adjacent pairs of pants at their endpoints, and normalising the resulting curves. Again, since the curves in \mathcal{A}_1 are short, these operations can be performed directly in polynomial time in ξ . Moreover, the multicurve A_2 will have $\mathcal{O}(\xi)$ components, each with complexity $\mathcal{O}(\xi \cdot \log \xi)$.

Finally, we let \mathcal{A}_2 be the set of components of A_2 , and set $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$. We invoke the Alexander method [12, Proposition 2.8] again to conclude that a homeomorphism $S \rightarrow S$ fixing the isotopy class of each curve in \mathcal{A} is periodic. \square

4.2 Subsurface projection

A key ingredient in Bell and Webb’s algorithm – and, therefore, in ours as well – is the bounded geodesic image theorem, as proved by Masur and Minsky [25] and later refined by Webb [34]. We quickly introduce the relevant notations and definitions, following the first reference [25, Sections 2.3 and 2.4].

A subspace $Y \subseteq S$ is an *essential surface* if it is a component of $S \setminus \Sigma$ for some essential multicurve $\Sigma \subseteq S$; note that we do not assume that $\xi(Y) \geq 2$. Suppose that Y




Figure 4.1. The multicurve A_2 intersects each pair of pants P in the pants decomposition induced by A_1 in three arcs, one for each pair of boundary components of P .

is not an annulus nor a sphere with 3 punctures. The *subsurface projection map* is a function

$$\pi_Y: \mathcal{C}(S) \longrightarrow \mathcal{PC}_0(Y),$$

where $\mathcal{PC}_0(Y)$ denotes the set of finite subsets of $\mathcal{C}_0(Y)$, satisfying the following properties for every $a \in \mathcal{C}_0(S)$:

- 1 $\pi_Y(a) = \emptyset$ if and only if a can be isotoped outside Y ;
- 2 $f(\pi_Y(a)) = \pi_{f(Y)}(f(a)) \in \mathcal{PC}_0(f(Y))$ for every homeomorphism $f: S \rightarrow S$.

We refer the reader to [25, Section 2.3] for the actual definition.

If Y is an annulus, we define $\mathcal{C}_0(Y)$ to be the set of arcs in the closure $\text{clos}(Y)$ of Y connecting the two boundary components of $\text{clos}(Y)$, modulo isotopies fixing ∂Y pointwise. The curve graph $\mathcal{C}(Y)$ is obtained by connecting two classes of arcs in $\mathcal{C}_0(S)$ by an edge of length 1 if they admit representatives with disjoint interiors. Let c be the core curve of Y . The subsurface projection map

$$\pi_Y: \mathcal{C}_0(S) \longrightarrow \mathcal{PC}_0(Y)$$

can be defined also in this case, and satisfies the following property: for every curve $a \in \mathcal{C}_0(S)$, the projection $\pi_Y(a)$ is empty if and only if $i(a, c) = 0$. Moreover, the following lemma shows how the subsurface projection map behaves nicely under homeomorphisms that are isotopic to Dehn twists in a sufficiently large neighbourhood of Y .

Lemma 4.3. *Let $Y \subseteq S$ be an essential annulus with core curve c , and $Z \subseteq S$ be an essential subsurface such that $Y \subseteq Z$ and c is an essential curve in Z . Let $f: S \rightarrow S$ be a homeomorphism preserving Z such that $f|_Z$, considered as a homeomorphism of Z , is isotopic to T_c^k for some integer $k \neq 0$. Then*

$$\text{dist}_{\mathcal{C}(Y)}(\pi_Y(a), \pi_Y(f(a))) \geq |k| - 1$$

for every curve $a \in \mathcal{C}_0(S)$ such that $\pi_Y(a) \neq \emptyset$.

Proof. We recall from [25, Section 2.4] the definition of π_Y in the annular case, which is not needed anywhere in this article except for the proof of this lemma. Let us fix a hyperbolic metric for S . Let $q: A \rightarrow S$ be the cover associated to the group $\pi_1 Y < \pi_1 S$. The surface A is topologically an open annulus, and can be compactified to a closed

annulus \bar{A} using the hyperbolic metric inherited from S . Then the subsurface projection map is the function

$$\pi_Y: \mathcal{C}_0(S) \longrightarrow \mathcal{PC}_0(A)$$

sending a curve $a \in \mathcal{C}_0(S)$ to the set of lifts of a to A whose closures in \bar{A} connect the two components of $\partial\bar{A}$. This is a finite set of diameter at most 1 in $\mathcal{C}(A)$, and is empty precisely when a can be isotoped away from c . We can think of π_Y as a map to $\mathcal{PC}_0(Y)$ if we fix an orientation-preserving homeomorphism $\bar{A} \rightarrow \text{clos}(Y)$.

Let $\bar{c} \subseteq A$ be the unique lift of c which is a core curve of \bar{A} , and let \bar{Y} and \bar{Z} be the components of $q^{-1}(Y)$ and $q^{-1}(Z)$ respectively containing \bar{c} . Finally, let $\bar{f}: \bar{A} \rightarrow \bar{A}$ be the lift of f that preserves \bar{c} . Note that, since c is essential in Z , the closure of \bar{Z} in \bar{A} intersects both boundary components of \bar{A} .

Up to isotopy, we can assume that f restricts to the identity on $Z \setminus Y$. It follows that \bar{f} is the identity on the two components of $\bar{Z} \setminus q^{-1}(Y)$ adjacent to \bar{Y} . In particular, the homeomorphism \bar{f} will be the identity on

$$W = \partial\bar{Y} \cup b_1 \cup b_2,$$

where $b_1, b_2 \subseteq \text{clos}(\bar{Z} \setminus q^{-1}(Y))$ are arcs connecting the two boundary components of \bar{Y} to $\partial\bar{A}$. See Figure 4.2 for a graphical representation of the objects introduced so far.




Figure 4.2. (a) Several lifts of c to \bar{A} , including the core curve \bar{c} . (b) The annulus $\bar{Y} \subseteq \bar{A}$. (c) The lift \bar{Z} of Z containing \bar{c} . (d) A subspace $W \subseteq \bar{A}$ on which \bar{f} is the identity.

Let \bar{a} be any arc in $\mathcal{C}_0(A)$. Note that $A \setminus W$ is the disjoint union of two open discs and the annulus \bar{Y} ; since \bar{f} is the identity on W and is the k -th power of the Dehn twist about \bar{c} on \bar{Y} , we can isotope $\bar{f}(\bar{a})$, fixing $\partial\bar{A}$, so that

- it does not intersect \bar{a} in $A \setminus \text{clos}(\bar{Y})$,
- it intersects \bar{a} twice in $\partial\bar{Y}$, and
- it intersects \bar{a} exactly $|k| - 1$ times in \bar{Y} , always with the same sign,




Figure 4.3. Isotoping $\bar{f}(\bar{a})$ (shown in (a)) so that it intersects \bar{a} at least $|k| - 1$ times (b).

as shown in Figure 4.3

It follows that the absolute value of the algebraic intersection number of \bar{a} and $\bar{f}(\bar{a})$ is bounded below by $|k| - 1$ (note that the two intersections on $\partial\bar{Y}$ might not be transverse). Equation (2.3) in [25] then implies that

$$\text{dist}_{\mathcal{C}(A)}(\bar{a}, \bar{f}(\bar{a})) \geq |k|.$$

Finally, we deduce the inequality in the statement of the lemma by recalling that $\pi_Y(a)$ has diameter at most 1 in $\mathcal{C}(A)$, and realising that $\pi_Y(f(a)) = \bar{f}(\pi_Y(a))$. \square

We can now state the bounded geodesic image theorem, combining Theorems 4.1.8 and 4.2.1 of Webb [34]’s thesis.

Theorem 4.4. *Let Y be an essential subsurface of S which is not a sphere with 3 punctures. Let $\gamma \subseteq \mathcal{C}(S)$ be a geodesic such that $\pi_Y(a) \neq \emptyset$ for every vertex $a \in \gamma \cap \mathcal{C}_0(S)$. Then*

$$\text{diam}_{\mathcal{C}(Y)}(\pi_Y(\gamma \cap \mathcal{C}_0(S))) \leq D_{\text{b.g.i}},$$

where

$$D_{\text{b.g.i}} = 100.$$

The reason why we take $D_{\text{b.g.i}}$ to be 100 is that we (following Masur and Minsky) define the subsurface projection map in the non-annular case to be the composition of Webb’s map $\kappa_Y : \mathcal{C}_0(S) \rightarrow \mathcal{PAC}_0(Y)$ to finite subsets of the *arc and curve graph* with the 2-Lipschitz function $\psi : \mathcal{AC}_0(Y) \rightarrow \mathcal{PC}_0(Y)$ introduced in [25, Lemma 2.2]. Therefore, Webb’s bound of 50 for non-annular subsurfaces must be doubled.

4.3 Stable translation length

The quantity we will use to distinguish reducible homeomorphisms from pseudo-Anosov ones is the stable translation length of their action on the curve graph. For a homeomorphism $f : S \rightarrow S$, define its *stable translation length* to be

$$\lambda(f) = \liminf_{k \rightarrow \infty} \frac{\text{dist}_{\mathcal{C}(S)}(a, f^k(a))}{k},$$

where a is an arbitrary curve in $\mathcal{C}_0(S)$, the choice of which is clearly immaterial. It is easy to see that, if f is periodic or reducible, then $\lambda(f) = 0$. The converse is also true,

as proved by Masur and Minsky [24, Proposition 4.6]. In fact, they show that if f is pseudo-Anosov then $\lambda(f) \geq \lambda_0$ for some constant $\lambda_0 > 0$ depending only on S . Later, Gadre and Tsai [14] compute an explicit lower bound for λ_0 , which we take as an equality for the purpose of this article:

$$\lambda_0 = \frac{1}{162\chi^2}.$$

We remark that the result by Gadre and Tsai is only stated for surfaces with $\xi(S) \geq 2$. However, Theorem 18 in [2] implies that the stable translation lengths of pseudo-Anosov homeomorphisms of the torus with one puncture and of the sphere with 4 punctures are integers. In particular, the lower bound for λ_0 stated above holds for surfaces S with $\xi(S) = 1$ as well.

When f is pseudo-Anosov, we have the following easy estimate.

Proposition 4.5. *Let S be a surface with $\xi(S) \geq 1$, and let $f: S \rightarrow S$ be a pseudo-Anosov homeomorphism. Then for every integer $h \geq 0$ and every curve $a \in \mathcal{C}_0(S)$ the following inequality holds:*

$$\text{dist}_{\mathcal{C}(S)}(a, f^h(a)) \geq \lambda_0 \cdot h.$$

Proof. Note that, for every integer $h_1 \geq 0$, we have

$$\text{dist}_{\mathcal{C}(S)}(a, f^{hh_1}(a)) \leq h_1 \text{dist}_{\mathcal{C}(S)}(a, f^h(a)).$$

As a consequence, we find that

$$\lambda_0 \leq \lambda(f) \leq \liminf_{h_1 \rightarrow \infty} \frac{\text{dist}_{\mathcal{C}(S)}(a, f^{hh_1}(a))}{h \cdot h_1} \leq \frac{\text{dist}_{\mathcal{C}(S)}(a, f^h(a))}{h}. \quad \square$$

Our goal is now to estimate how quickly the ratio $\text{dist}_{\mathcal{C}(S)}(a, f^k(a))/k$ converges to 0 as $k \rightarrow \infty$ when f is a non-periodic reducible homeomorphism. The following proposition is based on [5, Proposition 4.4], with some additional effort on our part to keep the constants polynomial in ξ .

Proposition 4.6. *Let $f: S \rightarrow S$ be a non-periodic reducible homeomorphism. Then there exists an integer $1 \leq h_0 \leq C_{\text{red}}$ such that for every integer $h \geq 0$ and every curve $a \in \mathcal{C}_0(S)$ the inequality*

$$\text{dist}_{\mathcal{C}(S)}(a, f^{hh_0}(a)) \leq \text{dist}_{\mathcal{C}(S)}(a, f^{h_0}(a)) + 2$$

holds. The constant C_{red} is given by

$$C_{\text{red}} = 2(D_{\text{b.g.i}} + 2) \cdot \lambda_0^{-1} \cdot \xi.$$

Proof. Let Σ be the *canonical reduction system* of f as defined by Handel and Thurston [16, §2] (there, the homeomorphism is called τ and the reduction system Γ). More explicitly, let Σ' be the set of all vertices $c \in \mathcal{C}_0(S)$ whose orbit under f is finite, and let Σ be the union of all the curves in Σ' that can be realised disjointly from all other curves in Σ' . Lemma 2.2 of [16] shows that Σ is non-empty, and moreover satisfies the following property: for every component Z of $S \setminus \Sigma$ and every integer $k \geq 1$ such that $f^k(Z) = Z$, the homeomorphism $f^k|_Z: Z \rightarrow Z$ is periodic or pseudo-Anosov.

Note that Σ has at most ξ components; let c be one of them. There is an integer $1 \leq h_1 \leq 2\xi$ such that the homeomorphism $f_1 = f^{h_1}$ preserves c and the two (not necessarily distinct) components Z_1 and Z_2 of $S \setminus \Sigma$ which are adjacent to c . Let

$f_2 = f_1^{h_2}$, where h_2 is a positive integer that will be defined shortly; its value will only depend on the quantities introduced so far. Let $a \in \mathcal{C}_0(S)$ be any curve, and fix a geodesic $\gamma \subseteq \mathcal{C}(S)$ between a and $f_2(a)$. We show that $\text{dist}_{\mathcal{C}(S)}(\gamma, c) \leq 1$ by analysing two cases.

□ Suppose that the restriction of f_1 to (say) Z_1 is pseudo-Anosov; in particular, this entails that $\xi(Z_1) \geq 1$. Let $h_2 = \lfloor D_{\text{b.g.i}} \cdot \lambda_0^{-1} \rfloor + 1$. Suppose, for a contradiction, that $\text{dist}_{\mathcal{C}(S)}(\gamma, c) \geq 2$. Then every vertex of γ must intersect Z_1 , and Theorem 4.4 gives the bound

$$\text{diam}_{\mathcal{C}(Z_1)}(\pi_{Z_1}(\gamma \cap \mathcal{C}_0(S))) \leq D_{\text{b.g.i}}.$$

Fix now a curve $\bar{a} \in \pi_{Z_1}(a)$. We can apply Proposition 4.5 to the surface Z_1 and the homeomorphism $f_1|_{Z_1}$ and deduce the chain of inequalities

$$\begin{aligned} \lambda_0(S) \cdot h_2 &\leq \lambda_0(Z_1) \cdot h_2 \\ &\leq \text{dist}_{\mathcal{C}(Z_1)}(\bar{a}, f_2(\bar{a})) \\ &\leq \text{diam}_{\mathcal{C}(Z_1)}(\pi_{Z_1}(a), f_2(\pi_{Z_1}(a))) \\ &= \text{diam}_{\mathcal{C}(Z_1)}(\pi_{Z_1}(a), \pi_{Z_1}(f_2(a))) \\ &\leq \text{diam}_{\mathcal{C}(Z_1)}(\pi_{Z_1}(\gamma \cap \mathcal{C}_0(S))) \leq D_{\text{b.g.i}}, \end{aligned}$$

from which we derive the contradiction $h_2 \leq D_{\text{b.g.i}} \cdot \lambda_0^{-1}$.

□ Suppose now that the restrictions of f_1 to Z_1 and Z_2 are periodic. Let $Z = Z_1 \cup Z_2 \cup c$, and let $Y \subseteq Z$ be a small annular neighbourhood of c . There exists an integer

$$1 \leq h_3 \leq C_{\text{per}}(Z_1) \cdot C_{\text{per}}(Z_2) \leq C_{\text{per}}(S)^2$$

such that $f_1^{h_3}|_{Z_1}$ and $f_1^{h_3}|_{Z_2}$ are isotopic to the identity. Note that $f_1^{h_3}|_Z$ is isotopic to T_c^k on Z for some $k \in \mathbb{Z}$; the integer k is non-zero, otherwise $f_1|_Z$ would be periodic and c would not be part of the canonical reduction system Σ .

Let $h_2 = (D_{\text{b.g.i}} + 2) \cdot h_3$, and suppose for a contradiction that $\text{dist}_{\mathcal{C}(S)}(\gamma, c) \geq 2$. Then every vertex of γ intersects Y , and again by Theorem 4.4 we get the bound

$$\text{diam}_{\mathcal{C}(Y)}(\pi_Y(\gamma \cap \mathcal{C}_0(S))) \leq D_{\text{b.g.i}}.$$

On the other hand, since c is clearly essential in Z , Lemma 4.3 applied to f_2 implies that

$$\text{dist}_{\mathcal{C}(Y)}(\pi_Y(a), \pi_Y(f_2(a))) \geq |(D_{\text{b.g.i}} + 2) \cdot k| - 1 \geq D_{\text{b.g.i}} + 1,$$

contradicting Theorem 4.4.

In both cases, we have shown the existence of an integer

$$1 \leq h_2 \leq (D_{\text{b.g.i}} + 2) \cdot \lambda_0^{-1}$$

such that $\text{dist}_{\mathcal{C}(S)}(\gamma, c) \leq 1$; let $b \in \gamma \cap \mathcal{C}_0(S)$ be at distance at most 1 from c . We set $h_0 = h_1 \cdot h_2$, noting that the bound $h_0 \leq C_{\text{red}}$ is satisfied. If now $h \geq 1$ is an arbitrary integer, we have the following chain of inequalities:

$$\begin{aligned} \text{dist}_{\mathcal{C}(S)}(a, f_2^h(a)) &\leq \text{dist}_{\mathcal{C}(S)}(a, c) + \text{dist}_{\mathcal{C}(S)}(c, f_2^h(a)) \\ &= \text{dist}_{\mathcal{C}(S)}(a, c) + \text{dist}_{\mathcal{C}(S)}(c, f_2(a)) \\ &\leq \text{dist}_{\mathcal{C}(S)}(a, b) + \text{dist}_{\mathcal{C}(S)}(b, f_2(a)) + 2 \text{dist}_{\mathcal{C}(S)}(b, c) \\ &\leq \text{dist}_{\mathcal{C}(S)}(a, f_2(a)) + 2. \end{aligned}$$

The equality on the second line follows from the fact that c is fixed by f_2 ; the last inequality holds because b lies on a geodesic between a and $f_2(a)$, and $\text{dist}_{\mathcal{C}(S)}(b, c) \leq 1$. This concludes the proof. □

4.4 Algorithmic classification

We are now ready to present our algorithm for the Nielsen-Thurston classification of surface homeomorphisms.

Theorem 4.7 (Nielsen-Thurston classification algorithm). *Let $f: S \rightarrow S$ be a homeomorphism. There is an algorithm to decide the Nielsen-Thurston type of f . The running time of the algorithm is polynomial in ξ and $\|f\|$.*

Proof. Proposition 4.2 gives a procedure to decide whether f is periodic. Assuming it is not, we now need to decide if f is reducible and non-periodic or pseudo-Anosov. To this aim, we first find a curve $a \in \mathcal{C}_0(S)$ such that $\mathcal{T}(a)_e \leq 2$ for every edge e of \mathcal{T} ; for instance, we can take a to be a component of a neighbourhood of an edge of \mathcal{T} . In particular, note that $\|a\|_1 \leq \log(-6\chi + 6)$. For every integer $1 \leq h_0 \leq C_{\text{red}}$, we perform the following operations.

- Compute

$$h = \left\lfloor \frac{L_{\times} \cdot (2\|f^{h_0}(a)\|_1 + L_+ + 6) + L_+}{\lambda_0 \cdot h_0} \right\rfloor + 1.$$

- Apply Proposition 3.9 to the curves a and $f^{hh_0}(a)$ to obtain an integer $d \geq 0$ such that

$$d - L_+ \leq \text{dist}_{\mathcal{C}(S)}(a, f^{hh_0}(a)) \leq L_{\times} \cdot d + L_+.$$

- Check whether the inequality

$$d \leq 2\|f^{h_0}(a)\|_1 + L_+ + 6 \tag{14}$$

is satisfied.

We claim that f is reducible if and only if (14) is satisfied for at least one value of $1 \leq h_0 \leq C_{\text{red}}$. Since the procedure described above can be performed in polynomial time in ξ and $\|f\|$, it is enough to prove this claim.

• If f is reducible (and non-periodic), then Proposition 4.6 guarantees the existence of an integer $1 \leq h_0 \leq C_{\text{red}}$ such that

$$\begin{aligned} d - L_+ &\leq \text{dist}_{\mathcal{C}(S)}(a, f^{hh_0}(a)) \\ &\leq \text{dist}_{\mathcal{C}(S)}(a, f^{h_0}(a)) + 2 \\ &\leq 2 \log i(a, f^{h_0}(a)) + 4 \\ &\leq 2\|f^{h_0}(a)\|_1 + 6, \end{aligned}$$

which obviously implies that (14) holds for this value of h_0 .

• If f is pseudo-Anosov and h_0 is any positive integer, then Proposition 4.5 gives

$$\begin{aligned} L_{\times} \cdot d + L_+ &\geq \text{dist}_{\mathcal{C}(S)}(a, f^{hh_0}(a)) \\ &\geq \lambda_0 \cdot h_0 \cdot h \\ &> L_{\times} \cdot (2\|f^{h_0}(a)\|_1 + L_+ + 6) + L_+, \end{aligned}$$

which contradicts (14). □

Remark 4.8. In the statement of Theorem 4.7, we give no estimate for the degree of the polynomial dependence of the running time on $\|f\|$. This omission stems from three main reasons: firstly, since we are not imposing any specific representation of surface homeomorphisms, we have no control over how expensive it is to compute images of curves under f ; secondly, we make use of Proposition 3.9, for which – as pointed out in Remark 3.10 – we do not give an explicit polynomial dependence in the closed case; thirdly, the proof of Proposition 4.2 for closed surfaces relies on the algorithm of [22, Theorem 1.2], for which Lackenby does not give an explicit polynomial bound on the running time.

However, with slight loss of generality, we can give a more precise upper bound to the running time of the algorithm of Theorem 4.7, assuming that S is punctured, and f is given as a sequence of $\|f\|$ flips (and no Dehn twists) in the flip and twist graph. In this setting, it is easy to see that, given an essential curve c on S , we can compute $f(c)$ in time $\mathcal{O}(\text{poly}(\xi) \cdot \|f\| \cdot \|c\|)$; moreover, the complexity of $f(c)$ satisfies the bound

$$\|f(c)\|_1 \leq \|f\| + \|c\|_1.$$

It follows that the integer h used in the proof of Theorem 4.7 satisfies $h = \mathcal{O}(\text{poly}(\xi) \cdot \|f\|)$, and hence

$$\|f^{h_0}(a)\|_1 = \mathcal{O}(\text{poly}(\xi) \cdot \|f\|^2).$$

Using the bound given in Remark 3.10 for the running time of Proposition 3.9 for punctured surfaces – and noting that the algorithm of Proposition 4.2 runs in linear time in $\|f\|$ when S is punctured – we conclude that, in this restricted setting, the Nielsen-Thurston classification algorithm has running time

$$\mathcal{O}(\text{poly}(\xi) \cdot \|f\|^4 \cdot \log\|f\|). \quad \times$$

Constants

For ease of navigation, we have collected all the constants introduced in this article in the following table.

Constant	Value	Page
g	genus	–
p	number of punctures	–
χ	Euler characteristic	–
ξ	$3g - 3 + p$	4
C_{per}	-6χ	37
C_{red}	$2(D_{\text{b.g.i}} + 2) \cdot \lambda_0^{-1} \cdot \xi$	42
D	$D_{\text{d.n}} + D_{\text{prox}} \cdot N_{\text{dec}}$	29
$D_{1.s}$	$48\xi^2 \cdot D_{\text{prox}}$	28
$D_{\text{b.g.i}}$	100	41

Continued on next page

(Continued)

Constant	Value	Page
$D_{d.n}$	$12\xi \cdot D'_{d.n}$	27
$D'_{d.n}$	$D_{\text{prox}} + 10$	25
D_{prox}	17	20
L_{\times}	$2D + 12$	29
L_{+}	$L_{+}^0 + 36$	36
L_{+}^0	$D_{1.s} + 2D + 8$	29
N_{dec}	$6\xi + 1$	23
λ_0	$\frac{1}{162\chi^2}$	42

References

- [1] Ian Agol, Joel Hass, and William Thurston. “3-manifold knot genus is NP-complete”. In: *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*. ACM, New York, 2002, pp. 761–766.
- [2] Hyungryul Baik, Changsub Kim, Sanghoon Kwak, and Hyunshik Shin. “On translation lengths of Anosov maps on the curve graph of the torus”. In: *Geom. Dedicata* 214 (2021), pp. 399–426.
- [3] Alan F. Beardon, Meira Hockman, and Ian Short. “Geodesic continued fractions”. In: *Michigan Math. J.* 61.1 (2012), pp. 133–150.
- [4] Marc C. Bell. “Recognising mapping classes”. PhD thesis. University of Warwick, 2015.
- [5] Marc C. Bell and Richard C. H. Webb. “Polynomial-time algorithms for the curve graph”. Preprint ([arXiv:1609.09392](https://arxiv.org/abs/1609.09392)). 2016.
- [6] Mark C. Bell. “Simplifying triangulations”. In: *Discrete Comput. Geom.* 66.1 (2021), pp. 1–11.
- [7] Mladen Bestvina and Michael Handel. “Train-tracks for surface homeomorphisms”. In: *Topology* 34.1 (1995), pp. 109–140.
- [8] Joan Birman, Dan Margalit, and William Menasco. “Efficient geodesics and an effective algorithm for distance in the complex of curves”. In: *Math. Ann.* 366.3-4 (2016), pp. 1253–1279.
- [9] Brian H. Bowditch. “Intersection numbers and the hyperbolicity of the curve complex”. In: *J. Reine Angew. Math.* 598 (2006), pp. 105–129.
- [10] Brian H. Bowditch. “Uniform hyperbolicity of the curve graphs”. In: *Pacific J. Math.* 269.2 (2014), pp. 269–280.
- [11] Matthieu Calvez. “Fast Nielsen-Thurston classification of braids”. In: *Algebr. Geom. Topol.* 14.3 (2014), pp. 1745–1758.
- [12] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Vol. 49. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 2012, pp. xiv+472.

- [13] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston’s work on surfaces*. Vol. 48. Mathematical Notes. Translated from the 1979 French original by Djun M. Kim and Dan Margalit. Princeton University Press, Princeton, NJ, 2012, pp. xvi+254.
- [14] Vaibhav Gadre and Chia-Yen Tsai. “Minimal pseudo-Anosov translation lengths on the complex of curves”. In: *Geom. Topol.* 15.3 (2011), pp. 1297–1312.
- [15] Hessam Hamidi-Tehrani and Zong-He Chen. “Surface diffeomorphisms via train-tracks”. In: *Topology Appl.* 73.2 (1996), pp. 141–167.
- [16] Michael Handel and William P. Thurston. “New proofs of some results of Nielsen”. In: *Adv. in Math.* 56.2 (1985), pp. 173–191.
- [17] David Harvey and Joris van der Hoeven. “Integer multiplication in time $O(n \log n)$ ”. In: *Ann. of Math. (2)* 193.2 (2021), pp. 563–617.
- [18] William J. Harvey. “Boundary structure of the modular group”. In: *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*. Vol. 97. Ann. of Math. Stud. Princeton Univ. Press, Princeton, N.J., 1981, pp. 245–251.
- [19] Allen Hatcher. “On triangulations of surfaces”. In: *Topology Appl.* 40.2 (1991), pp. 189–194.
- [20] John Hempel. “3-manifolds as viewed from the curve complex”. In: *Topology* 40.3 (2001), pp. 631–657.
- [21] Thomas Koberda and Johanna Mangahas. “An effective algebraic detection of the Nielsen-Thurston classification of mapping classes”. In: *J. Topol. Anal.* 7.1 (2015), pp. 1–21.
- [22] Marc Lackenby. “Some fast algorithms for curves in surfaces”. Preprint (arXiv:2401.16056). 2024.
- [23] Marc Lackenby and Mehdi Yazdi. “Bounds for the number of moves between pants decompositions, and between triangulations”. Preprint (arXiv:2401.14233). 2024.
- [24] Howard A. Masur and Yair N. Minsky. “Geometry of the complex of curves. I. Hyperbolicity”. In: *Invent. Math.* 138.1 (1999), pp. 103–149.
- [25] Howard A. Masur and Yair N. Minsky. “Geometry of the complex of curves. II. Hierarchical structure”. In: *Geom. Funct. Anal.* 10.4 (2000), pp. 902–974.
- [26] Howard A. Masur and Yair N. Minsky. “Quasiconvexity in the curve complex”. In: *In the tradition of Ahlfors and Bers, III*. Vol. 355. Contemp. Math. Amer. Math. Soc., Providence, RI, 2004, pp. 309–320.
- [27] Jakob Nielsen. “Abbildungsklassen endlicher Ordnung”. In: *Acta Math.* 75 (1943), pp. 23–115.
- [28] Robert C. Penner and John L. Harer. *Combinatorics of train tracks*. Vol. 125. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1992, pp. xii+216.
- [29] Marcus Schaefer, Eric Sedgwick, and Daniel Stefankovic. “Computing Dehn Twists and Geometric Intersection Numbers in Polynomial Time.” In: *Proceedings of the 20th Annual Canadian Conference on Computational Geometry, CCCG 2008*.
- [30] Saul Schleimer. *Notes on the complex of curves*. <http://homepages.warwick.ac.uk/~masgar/Maths/notes.pdf>.

- [31] Kenneth J. Shackleton. “Tightness and computing distances in the curve complex”. In: *Geom. Dedicata* 160 (2012), pp. 243–259.
- [32] Itaru Takarajima. “A combinatorial representation of curves using train tracks”. In: *Topology Appl.* 106.2 (2000), pp. 169–198.
- [33] Yohsuke Watanabe. “Intersection numbers in the curve complex via subsurface projections”. In: *J. Topol. Anal.* 9.3 (2017), pp. 419–439.
- [34] Richard C. H. Webb. “Effective geometry of curve graphs”. PhD thesis. University of Warwick, 2014.
- [35] Anders Wiman. “Über die hyperelliptischen Curven und diejenigen vom Geschlechte $p = 3$, welche eindeutigen Transformationen in sich zulassen”. In: *Bihang Kongl. Svenska Vetenskaps-Akademiens Handlingar* (1895-1896).