1

Notes on coarse median spaces

Brian H. Bowditch
Mathematics Institute
University of Warwick
Coventry CV4 7AL
United Kingdom

Abstract

These are notes from a mini-course lectured by Brian H. Bowditch on coarse median spaces given at *Beyond Hyperbolicity* in Cambridge in June 2016.

1.1 Introduction

These lecture notes give a brief summary of the notion of a "coarse median space" as defined in [Bo1] and motivated by the centroid construction given in [BM2]. The basic idea is to capture certain aspects of the large-scale "cubical" structure of various naturally-occurring spaces. Thus, a coarse median space is a geodesic metric space equipped with a ternary "coarse median" operation, defined up to bounded distance, and satisfying a couple of simple axioms. Roughly speaking, these require that any finite subset of the space can be embedded in a finite CAT(0) cube complex in such a way that the coarse median operation agrees, up to bounded distance, with the natural combinatorial median in such a complex. One could express everything in terms of CAT(0) cube complexes, but it is more convenient to formulate it in terms of median algebras (which are essentially equivalent structures for finite sets). One can apply this notion to finitely generated groups via their Cayley graphs. Examples of coarse median spaces include Gromov hyperbolic spaces, mapping class groups and Teichmüller spaces of compact surfaces, right-angled Artin groups and geometrically finite kleinian groups in any dimension. The notion is useful for establishing certain results such as coarse rank and quasi-isometric rigidity for such spaces.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

In Sections 1.2 and 1.3 we review some of the background to coarse geometry and to median algebras respectively. In Section 1.4 we combine these ideas to introduce the notion of a coarse median space. In Section 1.5 we discuss the geometry of the mapping class groups and Teichmüller spaces. In Section 1.6 we outline how the coarse median property is applied to such spaces via asymptotic cones.

I thank Benjamin Barrett for his excellent work in preparing these notes, based on a mini-course I gave at the meeting "Beyond Hyperbolicity" in Cambridge in June 2016. I take responsibility for any errors introduced by my subsequent editing and elaborations. I thank the organisers, Mark Hagen, Richard Webb and Henry Wilton, for their invitation to speak at the meeting.

1.2 Quasi-isometry invariants

We begin by making some basic definitions which describe the types of spaces we wish to discuss.

Let (X, ρ) be a metric space.

Definition 1.2.1 A *geodesic* in X is a path whose length is equal to the distance between its endpoints. We say that X is a *geodesic metric space* if every pair of points in X is the pair of endpoints of some geodesic.

All of the metric spaces of interest in this paper will be geodesic spaces (though we only make this hypothesis where we need it).

Definition 1.2.2 A geodesic space X is *proper* if it is complete and locally compact.

(This is equivalent to saying that all closed bounded subsets of X are compact.)

Definition 1.2.3 Let (X, ρ) and (X', ρ') be geodesic metric spaces. We say that a map $\phi\colon X \to X'$ is *coarsely-Lipschitz* if there exist constants $k_1, k_2 \geq 0$ such that

$$\rho'(\phi(x), \phi(y)) \leq k_1 \rho(x, y) + k_2$$

for any x and y in X.

We say that ϕ is a *quasi-isometric embedding* if it is coarsely-Lipschitz and there also exist constants $k_1', k_2' \geq 0$ such that

$$\rho(x, y) \leq k_1' \rho'(\phi(x), \phi(y)) + k_2'$$

for any x and y in X.

We say that ϕ is a *quasi-isometry* if it is a quasi-isometric embedding and there also exists a constant $k_3 \geq 0$ such that $X' = N(\phi(X), k_3)$; that is, X' is equal to the k_3-neighbourhood of the image of ϕ. In other words the image of ϕ is cobounded.

Note that in this definition we do not assume that the map ϕ is continuous.

Given geodesic spaces, X, Y, we write $X \preceq Y$ if there exists a quasi-isometric embedding $X \to Y$, and $X \sim Y$ if there exists a quasi-isometry $X \to Y$. Then the relations $\preceq$ and $\sim$ are both reflexive and transitive and $\sim$ is also symmetric. However $\preceq$ is not antisymmetric: there exist spaces X and Y such that $X \preceq Y$ and $Y \preceq X$ but $X \nsim Y$. For example, consider the following subsets of the euclidean plane, $\mathbb{R}^2$, given by

$$\begin{aligned}\{(x,y) \mid x,y \geq 0\} &\hookrightarrow \{(x,y) \mid (x \geq 0 \text{ and } y \geq 0) \text{ or } x = 0\} \\ &\hookrightarrow \{(x,y) \mid x \geq 0\} \\ &\sim \{(x,y) \mid x,y \geq 0\}\end{aligned}$$

in the induced path metrics. It is not hard to show that the intermediate spaces are not quasi-isometric to each other.

Example 1.2.4 For any $n \geq 1$, we have $[0,\infty)^n \sim \mathbb{R}^{n-1} \times [0,\infty)$. (Indeed, one can see easily that these spaces are bi-Lipschitz equivalent.) This half-space will appear again; we denote it H^n. Note that it is equipped with the restriction of the euclidean metric (not the hyperbolic metric).

Definition 1.2.5 Let a group Γ act on a proper geodesic metric space X by isometries. The action is *properly discontinuous* if for any compact subset K of X the set

$$\{g \in \Gamma \mid gK \cap K \neq \varnothing\}$$

is finite. (In this case the quotient space X/Γ is Hausdorff.)

The action is *cocompact* if X/Γ is compact.

When the action is cocompact, one can show that Γ must be finitely generated.

The geometry of a group is related to the geometry of the spaces on which it acts by the following theorem.

Theorem 1.2.6 (Švarc-Milnor) *If Γ acts on proper geodesic metric spaces X and X' properly discontinuously, cocompactly and by isometries, then $X \sim X'$. (Indeed, we can take the quasi-isometry to be equivariant.)*

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

Example 1.2.7 The action of a group Γ by left-translation on its Cayley graph $\Delta(\Gamma)$ with respect to any finite generating set is properly discontinuous and cocompact. It follows by Theorem 1.2.6 that any two such Cayley graphs for the same group are quasi-isometric.

Note: throughout this paper, unless otherwise stated, we assume that any connected graph is equipped with the combinatorial path metric, which assigns unit length to each edge.

Remark 1.2.8 We can often assume quasi-isometries to be continuous. For example, if $I \subset \mathbb{R}$ is an interval, then any quasi-isometric embedding $\phi: I \to X$ is within a bounded distance of a continuous map, and such a map is automatically also a quasi-isometric embedding. We refer to such a map as a *quasi-geodesic*.

Theorem 1.2.9 $\mathbb{R}^2 \nleq \mathbb{R}$.

Proof Suppose for contradiction that $\phi: \mathbb{R}^2 \to \mathbb{R}$ is a quasi-isometric embedding. Without loss of generality, ϕ is continuous (since a simple argument shows that it can always be approximated up to bounded distance by a continuous map). Let $S \subset \mathbb{R}^2$ be a round circle of large radius centred at the origin. By the Intermediate Value Theorem there exists x in S such that $\phi(x) = \phi(-x)$, which gives a contradiction, provided we choose the radius sufficiently large in relation to the quasi-isometric parameters. □

In fact the same argument (choosing the centre of the circle appropriately) shows that $H^2 \nleq \mathbb{R}$. Moreover, replacing the Intermediate Value Theorem with the Borsuk–Ulam theorem, one can see that $\mathbb{R}^{n+1} \nleq \mathbb{R}^n$ for any n, and therefore $\mathbb{R}^m \sim \mathbb{R}^n$ only when $m = n$. Indeed one can see that $H^{n+1} \nleq \mathbb{R}^n$. By related arguments one can also show that any quasi-isometric embedding of $\mathbb{R}^n$ into itself in necessarily a quasi-isometry.

Definition 1.2.10 If X is a geodesic space, the *euclidean rank* of X E-rk$(X) \in \mathbb{N} \cup \{\infty\}$ is defined to be the maximum n such that $\mathbb{R}^n \leq X$. The *half-space rank* of X, H-rk(X), is defined to be the maximum n such that $H^n \leq X$.

Clearly, H-rk$(X) - 1 \leq$ E-rk$(X) \leq$ H-rk(X). These ranks are quasi-isometry invariants.

Note that, by the above observations, we have E-rk$(\mathbb{R}^n) =$ H-rk$(\mathbb{R}^n) = n$ and E-rk$(H^n) + 1 =$ H-rk$(H^n) = n$.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

Definition 1.2.11 A map $f\colon [0,\infty) \to [0,\infty)$ is an *isoperimetric bound* for X if there exists a constant k such that if $\gamma\colon S^1 \to X$ is any curve, we can cut γ into at most $f(\text{length}(\gamma))$ loops of length at most k.

(More formally, we can extend f to a map of the 1-skeleton of a cellulation of the disc, with boundary S^1, such that the length of the f-image of the boundary of any 2-cell has length at most k.)

The rate of growth of the isoperimetric bound is a quasi-isometry invariant. (Here the "growth rate" is interpreted up to linear bounds: we allow for linear reparametrisation of the domain and range of f.) In particular, we can talk about spaces with linear, quadratic and exponential isoperimetric bounds, et cetera.

A central notion in the subject is that of *Gromov hyperbolicity* [G1]. There are numerous equivalent definitions, among which we choose the following.

Definition 1.2.12 A geodesic metric space X is *hyperbolic* if there exists a constant k such that for any geodesic triangle in X, there exists a point m in X within distance k of each of the three sides of the triangle. (A "geodesic triangle" consists of three geodesic segments — its "sides" — cyclically connecting three points.)

It turns out that, up to bounded distance, m depends only on the vertices of the triangle, so if x, y and z are the vertices then we write $m = m(x,y,z)$.

This definition is quasi-isometry invariant. Moreover, Gromov showed that X is hyperbolic if and only if it has a linear isoperimetric bound. We note also the following geometric properties of hyperbolic spaces.

1 Hyperbolic metric spaces satisfy a *Morse Lemma*: any quasi-geodesic is close to any geodesic joining its end points. More precisely, the Hausdorff distance between them depends only on the quasi-isometry constants and the hyperbolicity constant k.
2 Hyperbolic metric spaces can be well approximated by trees: there exists a function $h\colon \mathbb{N} \to \mathbb{N}$ such that if X is k-hyperbolic and $A \subset X$ is a finite subset of cardinality at most p, there exists a tree $\tau \subset X$ with $A \subset \tau$ such that for any x and y in A, $\rho_\tau(x,y) \le \rho(x,y) + kh(p)$. Here ρ_τ denotes the induced path-metric on τ. (In this case we are allowing the edges of τ to have differing lengths.) Note that, using the Morse Lemma, it follows that the arc in τ from x to y is a bounded Hausdorff distance from any geodesic in X from x to y.

Also note that if X is hyperbolic, then H-rk$(X) \leq 1$.

Definition 1.2.13 Let a group Γ act on a geodesic space X by isometries. We say that the action is *quasi-isometrically rigid* if for any quasi-isometry $\phi\colon X \to X$ there exists $g \in \Gamma$ such that $\rho(gx, \phi(x)) \leq C$ for some constant C depending only on the quasi-isometry constants of the map.

When the group Γ is understood, we will express this by saying that X is "quasi-isometrically rigid".

1.3 Medians

We describe the basic properties of a median algebra and how they relate to CAT(0) cube complexes. Some basic references for median algebras are [BaH, Ro, Ve]. Some further discussion, relevant to these notes, is given in [Bo1, Bo4]. CAT(0) complexes are discussed, for example, in [BrH]. We can view a CAT(0) complex combinatorially as a simply-connected complex built out of cubes such that the link of every vertex is a flag simplicial complex. They are usually equipped with a euclidean (CAT(0)) cubical structure, though it is more natural to consider the ℓ^1 metric in the present context.

Let M be a set and let $\mu\colon M^3 \to M$ be a ternary operation. (Intuitively, we think of μ as mapping points a, b and c in M to a point "between a, b and c".)

The standard definition of a median algebra is simple, but somewhat formal and perhaps unintuitive.

Definition 1.3.1 (M, μ) is a *median algebra* if for any a, b, c, d and e in M,

(M1) $\mu(a, b, c) = \mu(b, a, c) = \mu(b, c, a)$,
(M2) $\mu(a, a, b) = a$ and
(M3) $\mu(a, b, \mu(c, d, e)) = \mu(\mu(a, b, c), \mu(a, b, d), e)$.

Given a and b in M we write $[a, b]_\mu = \{x \in M \mid \mu(a, b, x) = x\}$, which we abbreviate to $[a, b]$ if the choice of function μ is clear from context. The set $[a, b]$ is called the *interval* between a and b.

The notion of a median algebra can equivalently, and perhaps more intuitively, be formulated in terms of intervals. This follows from work of Sholander [Sho]. (See [Bo4] for some elaboration.)

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

Lemma 1.3.2 *Let M be a median algebra. The interval operation $[\cdot,\cdot]$ satisfies the following properties for any a, b, c in M:*

(I1) $[a,a] = \{a\}$,
(I2) $[a,b] = [b,a]$,
(I3) $c \in [a,b] \implies [a,c] \subset [a,b]$, *and*
(I4) there exists d (depending on a, b and c) such that $[a,b] \cap [b,c] \cap [c,a] = \{d\}$.

In property (I4) we can set $d = \mu(a,b,c)$.

We can alternatively view properties (I1)–(I4) as axioms, and we have the following converse for any set M.

Theorem 1.3.3 *[Sho] Given a map $[\cdot,\cdot]$ from M^2 to the power set $\mathcal{P}(M)$ satisfying axioms (I1)–(I4) above, there exists a map $\mu\colon M^3 \to M$ such that (M,μ) is a median algebra and $[\cdot,\cdot] = [\cdot,\cdot]_\mu$. In fact, we can set $\mu(a,b,c)$ to be the element d given in axiom (I4).*

Example 1.3.4 We give some examples of median algebras.

1 Let M be the two-point set $\{0,1\}$. Then there is a unique median algebra structure on M given by $\mu(0,0,0) = 0$, $\mu(0,0,1) = 0$, $\mu(0,1,1) = 1$, $\mu(1,1,1) = 1$ etc. (In other words μ represents the "majority vote".)
2 If M_1 and M_2 are median algebras then so is $M_1 \times M_2$, with the median defined separately on each co-ordinate.
3 Combining the previous two examples, the "n-cube" $\{0,1\}^n$ has a natural median algebra structure. One can show that any finite median algebra is a subalgebra of such a cube.
4 Trees are median algebras. Define the median of three points to be the centre of the tripod spanned by those points. Here a "tree" can be interpreted as a simplicial tree, or more generally any $\mathbb{R}$-tree. This includes the case of $\mathbb{R}$ itself: here the median of three points is just the point that lies between the other two.
5 Given any set X define a median on its power set $\mathcal{P}(X)$ by:

$$\begin{aligned}\mu(A,B,C) &= (A \cup B) \cap (B \cup C) \cap (C \cup A)\\ &= (A \cap B) \cup (B \cap C) \cup (C \cap A)\end{aligned}$$

for $A,B,C \subset X$. Then $(\mathcal{P}(X),\mu)$ is a median algebra.
6 The previous example generalises to any distributive lattice, with the median defined by a similar formula, using meets and joins in place of intersections and unions.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

7 Let Δ be a CAT(0) cube complex. Its vertex set $V(\Delta)$ can be made into a median algebra as follows. Let ρ be the combinatorial path metric on the 1-skeleton of Δ. Then given $a, b \in V(\Delta)$ let $[a,b]_\rho = \{x \in M : \rho(a,b) = \rho(a,x) + \rho(x,b)\}$. This definition satisfies axioms (I1)–(I4) above, so by Theorem 1.3.3 there exists a median algebra structure $\mu : V(\Delta)^3 \to V(\Delta)$ such that $[a,b]_\mu = [a,b]_\rho$.
8 $\mathbb{R}^n$ with the ℓ^1 metric, ρ. Here one defines the median similarly as in the previous example. This is median-isomorphic to the direct product of n copies of $\mathbb{R}$.
9 Similarly, CAT(0) cube complexes with the ℓ^1 metric (that is the path-metric obtained by putting the ℓ^1 metric on each cube). In this case, the vertex set is a subalgebra (that is, closed under μ).
10 More generally, a median metric space: that is any metric space (X, ρ) such that $[a,b]_\rho \cap [b,c]_\rho \cap [c,a]_\rho$ is a singleton for all $a, b, c \in X$ (which gives us the median of a, b, c). Note that this is just axiom (I4) in Theorem 1.3.3. Axioms (I1)–(I3) follow immediately from the metric space axioms.

A subset B of a median algebra M is a *subalgebra* if it is closed under μ. We write $B \leq M$. For any $A \subset M$, $\langle A \rangle \leq M$ is the subalgebra generated by A; that is, the intersection of all subalgebras of M containing A.

We say that a subset $C \subset M$ is *convex* if $[a,b] \subset C$ whenever $a, b \in C$. We note that convex sets are subalgebras, and that intervals themselves are convex.

The following are two basic facts about median algebras.

Theorem 1.3.5

1 Let M be a median algebra, and let $A \subset M$ with $|A| \leq p < \infty$. Then $|\langle A \rangle| \leq 2^{2^p}$.
2 Any finite median algebra is canonically the vertex set of a CAT(0) *cube complex.*

Note that these give rise to a third equivalent way of defining a median algebra: it is a set equipped with a ternary operation such that any finite subset is contained in another finite subset, closed under this operation, and isomorphic to the median structure on a finite CAT(0) cube complex.

In particular, in dealing with any finite subset of a median algebra, we can often just pretend we are living in a CAT(0) cube complex.

Definition 1.3.6 Define the *median rank* of M, M-rk(M), to be the maximum n such that $\{0,1\}^n \leq M$, so M-rk$(M) \in \mathbb{N} \cup \{\infty\}$.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

For example, M-rk$(\mathbb{R}^n) = n$, and if Δ is a CAT(0) cube complex then one can check that M-rk$(\Delta) =$ M-rk$(V(\Delta)) = \dim(\Delta)$ (that is, the standard notion of dimension — the maximal dimension of a cubical cell).

We now state two theorems about median metric spaces which will be useful for the discussion in Section 1.6.

Theorem 1.3.7 *[Bo1] Let M be a connected, locally convex topological median algebra of rank at most $n < \infty$. Then the locally compact dimension of M (i.e. the maximum topological dimension of a locally compact subset of M) is equal to the median rank* M-rk(M).

(For locally compact spaces, all of the standard definitions of topological dimension are equivalent [E]. For definiteness, we could take to mean covering dimension.)

Here a "topological median algebra" is simply one equipped with a topology with respect to which the median operation is continuous. It is "locally convex" if every point has a base of convex neighbourhoods. This is satisfied in the cases of interest here. For example, the hypotheses of the theorem hold in any finite-rank connected median metric space (as defined in Example 1.3.4 (10)).

Theorem 1.3.8 ([Bo4]) *If (M, ρ) is a connected complete finite-rank median metric space then there exists a canonical bi-Lipschitz-equivalent metric σ_ρ on M such that (M, σ_ρ) is* CAT(0).

We remark that, under the same hypotheses, one can also put a canonical bi-Lipschitz equivalent injective metric on M [Mi, Bo8].

In particular, it follows (from either the CAT(0) or injective metric) that M is contractible.

1.4 Coarse median spaces

Let (Λ, ρ) be a geodesic space.

Definition 1.4.1 A map $\mu\colon \Lambda^3 \to \Lambda$ is a *coarse median* if

(C1) There exist k and l such that for any $a, b, c, a', b', c' \in M$,

$$\rho\left(\mu(a,b,c), \mu(a',b',c')\right) \le k\left(\rho(a,a') + \rho(b,b') + \rho(c,c')\right) + l.$$

(C2) There exists $h\colon \mathbb{N} \to [0, \infty)$ such that if A is a subset of Λ containing

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

at most p points, then there exists a finite median algebra (Π, μ_Π) and maps

$$A \xrightarrow{\pi} \Pi \xrightarrow{\lambda} \Lambda$$

such that $\rho(a, \lambda\pi a) \le h(p)$ for all $a \in A$, and

$$\rho\left(\mu(\lambda a, \lambda b, \lambda c), \lambda\mu_\Pi(a, b, c)\right) \le h(p)$$

for all a, b and c in Π.

We say that Λ has *coarse (median) rank* at most n if (given some fixed map h) we can always take $\text{M-rk}(\Pi) \le n$. The coarse rank (i.e. the minimal such n) is denoted $\text{C-rk}(\Lambda)$.

Less formally, (C1) says that the coarse median is coarsely-Lipschitz, and (C2) says that, on finite sets, it looks like the median on a finite CAT(0) cube complex up to bounded distance. We remark that an equivalent set of axioms for a coarse median space has recently been described in [NWZ].

The existence of a coarse median on a geodesic space is a quasi-isometry invariant. We say that a finitely generated group is *coarse median* if some (hence any) Cayley graph with respect to a finite generating set admits a coarse median.

We give some examples.

Example 1.4.2

1 CAT(0) cube complexes (so right-angled Artin groups are coarse median groups).
2 Hyperbolic spaces: these spaces have coarse rank at most 1. (This follows from approximation of the space by a tree). In fact any coarse median space with coarse rank at most 1 is hyperbolic [Bo1] (see [NWZ] for a more direct proof).
3 The property is closed under taking direct products and relative hyperbolicity [Bo2].
4 From this it follows that any geometrically finite kleinian group (in any dimension) is coarse median. So are limit groups, as defined by Sela.
5 Mapping class groups, Teichmüller space in either the Teichmüller metric or the Weil–Petersson metric [Bo6, Bo5, Bo7], and the separating curve graphs [Vo]. (See Section 1.5).
6 Any hierarchically hyperbolic space [BHS1, BHS2].

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

The following is fairly easy to see, and only requires axiom (C1) [Bo1].

Theorem 1.4.3 *Any coarse median space satisfies a quadratic isoperimetric bound.*

One can also show the following [Bo1].

Theorem 1.4.4 *If* Λ *is a coarse median space, then* $\text{H-rk}(\Lambda) \leq \text{M-rk}(\Lambda)$.

We will outline how this is proven in Section 1.6. We will first elaborate on its consequences for the mapping class groups and Teichmüller space in the next section.

1.5 Surfaces

Let Σ be a compact orientable surface. Let $g(\Sigma)$ be its genus, $p(\Sigma)$ be the number of boundary components of Σ, and define the *complexity* of Σ to be

$$\xi(\Sigma) = 3g(\Sigma) - 3 + p(\Sigma).$$

This is the maximum number of disjoint curves (i.e. essential non-peripheral simple closed curves up to homotopy) that one can embed in Σ. We usually assume that $\xi(\Sigma) \geq 2$. We will denote the topological type of a surface of genus g and p boundary components by $S_{g,p}$.

Recall that the mapping class group, $\mathrm{Map}(\Sigma)$, can be defined as the group of self-homeomorphisms of Σ defined up to homotopy (or, equivalently, isotopy). This is a finitely presented group. For future reference, we note that $\mathbb{Z}^\xi \leq \mathrm{Map}(\Sigma)$. For example take the subgroup generated by Dehn twists around any maximal collection of disjoint simple closed curves (that is a "pants decomposition" of Σ).

We will focus on four particular spaces on which the mapping class group acts, namely, the marking graph $\mathcal{M}$, the curve graph $\mathcal{C}$, and Teichmüller space with the Teichmüller or Weil–Petersson metric, respectively denoted $\mathcal{T}$ and $\mathcal{W}$.

These spaces are interrelated. In fact, there are coarsely-Lipschitz $\mathrm{Map}(\Sigma)$-equivariant maps

$$\mathcal{M} \longrightarrow \mathcal{T} \longrightarrow \mathcal{W} \longrightarrow \mathcal{C}.$$

natural up to bounded distance.

We proceed to describe these spaces in more detail. We begin with

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

the marking graph. We write $\iota(\alpha, \beta)$ for the intersection number of two curves α, β (that is the minimal cardinality of $|\alpha \cap \beta|$ among realisations in Σ).

Definition 1.5.1 A set, a, of curves in Σ is said to *fill* Σ if for any curve γ in Σ, there is some $\alpha \in a$ with $\iota(\gamma, \alpha) \neq 0$. (Less formally, this says that a cuts Σ into discs and peripheral annuli.) A *marking* is a set, a, of curves in Σ that fills Σ and such that for all $\alpha, \beta \in a$, $\iota(\alpha, \beta) \leq 100$. The *marking graph*, $\mathcal{M}$, has vertex set $V(\mathcal{M})$ equal to the set of markings of Σ, and where two markings a and b in $\mathcal{M}$ are deemed adjacent if for any $\alpha \in a$ and $\beta \in b$ we have $\iota(\alpha, \beta) \leq 10000$.

(Here “100” and “10000” could be interpreted to mean any two sufficiently large numbers.)

The graph $\mathcal{M}$ is connected and $\mathrm{Map}(\Sigma)$ acts on $\mathcal{M}$ properly discontinuously and cocompactly. In particular, by Theorem 1.2.6 we see that $\mathcal{M}$ is quasi-isometric to (any Cayley graph of) $\mathrm{Map}(\Sigma)$. (A different definition is given in [MM2]. The notion is quite robust — any two sensible definitions will give equivariantly quasi-isometric graphs. It will not matter to us here which variation is chosen.)

One can show that the subgroup $\mathbb{Z}^\xi$ of $\mathrm{Map}(\Sigma)$ generated by Dehn twists is quasi-isometrically embedded. (This means that any orbit of this group in $\mathcal{M}$ is quasi-isometrically embedded.) From this we see that $\mathbb{R}^\xi \leq \mathcal{M}$. In other words, $\text{E-rk}(\mathcal{M}) \geq \xi$.

By a *Dehn twist flat* in $\mathcal{M}$ we will mean an $\mathbb{Z}^\xi$-orbit of this type, where the orbits are chosen to be uniformly quasi-isometrically embedded. Uniformity is possible since there are only finitely many conjugacy classes of subgroups of this type in $\mathrm{Map}(\Sigma)$. (Not all quasi-isometric embeddings of $\mathbb{R}^\xi$ into $\mathcal{M}$ arise from Dehn twist flats, however.)

Definition 1.5.2 We define the *curve graph*, $\mathcal{C}$, of Σ. Its vertex set $V(\mathcal{C})$ is the set of curves on Σ; two curves are deemed to be adjacent if they can be homotoped to be disjoint.

The curve graph is connected whenever $\xi(\Sigma) \geq 2$. In fact, the following result is central to the whole subject.

Theorem 1.5.3 *(Masur–Minsky) [MM1] $\mathcal{C}$ is hyperbolic.*

Recall that the *Teichmüller space* of Σ is the space of marked finite-area hyperbolic structures on $\Sigma - \partial\Sigma$, defined up to isotopy. (See, for example, [IT].) As a topological space it is homeomorphic to $\mathbb{R}^{2\xi}$, though we are interested here in its (large-scale) geometry. It admits many interesting

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

metrics; for example, the Teichmüller metric and the Weil–Petersson metric as mentioned above.

The spaces $\mathcal{T}$ and $\mathcal{W}$ have a somewhat different structure. Notably, $\mathcal{T}$ is complete, whereas $\mathcal{W}$ is not. The basic reason behind this can be thought of as follows. Take any essential non-peripheral simple closed curve on the surface. One can form a path in Teichmüller space by shrinking the length of the curve, while keeping the hyperbolic structure on the remainder of the surface approximately constant. In this way the surface develops an annular "Margulis tube", with our curve as its core. As the length of this curve tends to 0, the length of the tubes tends to ∞. We thus get a properly embedded path in Teichmüller space. In $\mathcal{T}$, this process takes an infinite amount of effort, and the path has infinite length. However, in $\mathcal{W}$ only a finite amount of effort is needed to pull the surface apart in this way, and the path has finite length. (See [W].)

In fact, one can take a maximal collection of disjoint curves in Σ and shrink these independently of each other. Since there are ξ such curves, this gives a proper map of $[0,\infty)^\xi$ into Teichmüller space. One can show that the map $[0,\infty) \to \mathcal{T}$ is a quasi-isometric embedding. Since $[0,\infty)^\xi \sim H^\xi$, we see that $H^\xi \le \mathcal{T}$. In other words, $\text{H-rk}(\mathcal{T}) \ge \xi$.

In $\mathcal{W}$ however, the image of this map is bounded, so we don't achieve very much by this. Instead, write

$$\xi_0 = \xi_0(\Sigma) = \lfloor (\xi(\Sigma)+1)/2 \rfloor.$$

The significance of this number is that we can cut Σ into ξ_0 pieces, each of complexity at least 1 (that means each contains an $S_{0,4}$ or an $S_{1,1}$). One can now deform the hyperbolic structures on these pieces independently, and by taking an appropriate bi-infinite path of such deformations in each component, we get a proper map of $\mathbb{R}^{\xi_0}$ into $\mathcal{W}$. One can also show that such a map is a quasi-isometric embedding. Therefore, $\text{E-rk}(\mathcal{W}) \ge \xi_0$.

In fact, one has equality when $\xi(\Sigma) \ge 2$, as the following theorem clarifies.

Theorem 1.5.4

1 *(Behrstock, Minsky, Hamenstädt) [BM1, H]* $\text{H-rk}(\mathcal{M}) = \text{E-rk}(\mathcal{M}) = \xi$.
2 *(Eskin, Masur, Rafi) [EMR1]* $\text{H-rk}(\mathcal{T}) = \xi$.
3 *(Eskin, Masur, Rafi) [EMR1]* $\text{H-rk}(\mathcal{W}) = \text{E-rk}(\mathcal{W}) = \xi_0$.

The remaining issue regarding $\text{E-rk}(\mathcal{T})$ is resolved by the following.

Theorem 1.5.5 *[Bo5]* $\mathbb{R}^\xi \le \mathcal{T}$ *if and only if* $g \le 1$ *or* $\Sigma = S_{2,0}$.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

We will outline later how Theorem 1.5.4 can also be derived from the coarse median property.

Recall that $\mathrm{Map}(\Sigma)$ acts naturally on $\mathcal{M}$, $\mathcal{T}$, $\mathcal{W}$ and $\mathcal{C}$. We briefly review some rigidity results for these actions.

Theorem 1.5.6

1 (Behrstock, Kleiner, Minsky, Mosher, Hamenstädt) [BKMM, H] $\mathcal{M}$ is quasi-isometrically rigid.

2 (Eskin, Masur, Rafi, Bowditch) [EMR2, Bo5] $\mathcal{T}$ is quasi-isometrically rigid.

3 (Bowditch) [Bo7] $\mathcal{W}$ is quasi-isometrically rigid if $g(\Sigma) + p(\Sigma) \geq 7$.

4 (Rafi, Schleimer) [RS] $\mathcal{C}$ is quasi-isometrically rigid.

It is a relatively simple matter to account for the low complexity cases ($\xi \leq 1$), so this give a compete answer for $\mathcal{M}$, $\mathcal{T}$ and $\mathcal{C}$. However, [Bo7] leaves unresolved about a dozen cases for $\mathcal{W}$.

We also have the following analogue of the cohopfian property (*cf.* the case of $\mathbb{R}^n$ discussed in Section 1.2).

Theorem 1.5.7 *[Bo6] Any quasi-isometric embedding of $\mathcal{M}$ into itself is a quasi-isometry.*

In fact, this is achieved by giving another proof of quasi-isometric rigidity of $\mathcal{M}$, but only using the weaker hypothesis that our map is a quasi-isometric embedding. (I do not know whether a similar statement holds for any of the other cases: $\mathcal{T}$, $\mathcal{W}$ or $\mathcal{C}$.)

It is time to explain how the coarse median property is brought into play.

Theorem 1.5.8 *[Bo6, Bo5, Bo7]. $\mathcal{M}$ and $\mathcal{T}$ are coarse median of rank ξ, $\mathcal{W}$ is coarse median of rank ξ_0 and $\mathcal{C}$ is coarse median of rank 1.*

(Of course, the last statement about $\mathcal{C}$ does not tell us anything essentially new — it follows directly from Theorem 1.5.3.)

From Theorem 1.4.3 it follows that each of these spaces satisfies a quadratic isoperimetric bound. So, for example, we recover the fact, due to Mosher [Mo], that the mapping class group has quadratic Dehn function. For $\mathcal{W}$ and $\mathcal{C}$ this follows respectively from the CAT(0) property and hyperbolicity of these spaces. The fact this holds for the Teichmüller metric appears to be new, though an independent proof has been announced by Kapovich and Rafi [KR].

Note that one can now also recover Theorem 1.5.4 using Theorem 1.4.4. (Another proof for $\mathcal{T}$ is given in [Du2], and for $\mathcal{T}$ and $\mathcal{W}$ in [BHS1].)

Theorem 1.5.5 still requires some more work, which we will not describe here.

As for the rigidity results, there is still quite a bit more to be done, but we will briefly discuss some of the ingredients in Section 1.6.

We spend the remainder of this section briefly describing how the coarse median structure arises in these situations. First, we consider the case of $\mathcal{M}$. This is based on the centroid construction of [BM2], which uses the notion of subsurface projection of Masur and Minsky [MM2].

Let $\mathcal{S}$ be the set of π_1-injective subsurfaces of Σ up to homotopy. For technical reasons it is helpful to exclude surfaces homeomorphic to $S_{0,1}$ or $S_{0,3}$.

For $\Phi \in \mathcal{S}$ one can define a coarsely-Lipschitz map $\theta_\Phi : \mathcal{M}(\Sigma) \to \mathcal{C}(\Phi)$. This definition is due to Masur and Minsky. We realise Φ and α to minimise the number of components of their intersection. For $a \in \mathcal{M}$ pick $\alpha \in a$ and let $\delta \subset \alpha$ be a component of $\alpha \cap \Phi$. Choose a curve γ in Φ with $\gamma \cap \delta = \varnothing$. Then let $\theta_\Phi(a) = \gamma$; this is well defined up to a bounded distance in $\mathcal{C}$. (Note that this definition assumes that Φ is not an annulus. In the case of an annulus a different definition is required. In this case, "$\mathcal{C}(\Phi)$" is a space quasi-isometric to the real line — in particular, hyperbolic. Only the logical structure is relevant here, so we will not elaborate on this point.) For a and b in $\mathcal{M}(\Sigma)$, define $\sigma_\Phi(a,b)$ to be the distance in $\mathcal{C}(\Phi)$ from $\theta_\Phi(a)$ to $\theta_\Phi(b)$.

The following theorem follows from work of Masur and Minsky.

Theorem 1.5.9 *[MM2] Let ρ be the distance function in $\mathcal{M}$. Then $\rho(a,b)$ is bounded above in terms of* $\max\{\sigma_\Phi(a,b) : \Phi \in \mathcal{S}\}$.

Theorem 1.5.10 *[BM2] For all $a, b, c, \in \mathcal{M}$ there exists $d \in \mathcal{M}$ such that for any $\Phi \in \mathcal{S}$,*

$$\rho\left(\theta_\Phi d, \mu_{\mathcal{C}(\Phi)}(\theta_\Phi a, \theta_\Phi b, \theta_\Phi c)\right)$$

is bounded by some constant depending only on the topological type of Σ.

Using Theorem 1.5.9 we see that d is well-defined up to bounded distance. Write $\mu(a,b,c) = d$ to get a map $\mathcal{M}^3 \to \mathcal{M}$. Using certain properties of subsurface projection one shows that this defines a coarse median structure on $\mathcal{M}$ [Bo1].

To obtain the bound on the rank of $\mathcal{M}$, one shows that if a quasi-square (that is, the image of a 2-cube under a coarse median homomorphism) has

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

a large projection to both Φ and Ψ in $\mathcal{S}$, then Φ and Ψ are either disjoint or equal. It then follows that $\text{M-rk}(M) \le \xi$ since this is the maximal number of disjoint elements of $\mathcal{S}$ that one can embed in Σ.

Similar constructions can be made to work for $\mathcal{T}$ and $\mathcal{W}$. For this one uses combinatorial models for these spaces. Specifically, it was shown by Brock [Br] that $\mathcal{W}$ is quasi-isometric to the so-called pants graph, and in [Ra] and [Du1] it is shown that $\mathcal{T}$ is quasi-isometric to the augmented marking graph. We will not give definitions here; suffice it to note that this allows us to employ similar arguments of subsurface projection. The key properties of subsurface projection needed are listed in [Bo6]. (See also [BHS2].)

We also remark that these models can be used to define the maps $\mathcal{M} \to \mathcal{T} \to \mathcal{W} \to \mathcal{C}$ mentioned earlier. (The composition of these maps, $\mathcal{M} \to \mathcal{C}$, simply selects one curve from the marking of the surface.)

The separating curve graph can be included in this picture as intermediate between $\mathcal{W}$ and $\mathcal{C}$. In most cases, it is coarse median of rank 2 [Vo]. As far as I know, its quasi-isometric rigidity has not been investigated.

1.6 Asymptotic cones

The rigidity of $\mathcal{M}$ is proven using a limiting argument phrased in terms of asymptotic cones (see [vdDW, G2]). We outline some of the ingredients here. We begin by defining the asymptotic cone of a metric space.

Let I be a countable set. Let $\mathcal{P} = \mathcal{P}(I)$ be the power set of I.

Definition 1.6.1 A subset $\mathcal{F} \subset \mathcal{P}$ is an *ultrafilter* if the following hold.

1 If $A, B \in \mathcal{F}$ then $A \cap B \in \mathcal{F}$.
2 If $A \in \mathcal{F}$ and $A \subset B$ then $B \in \mathcal{F}$.
3 If $A \subset I$ then either A or $I - A$ is in $\mathcal{F}$.
4 $\varnothing \notin \mathcal{F}$.

For example, if $a \in I$ then $\{A \in \mathcal{P}(I) \mid a \in A\}$ is an ultrafilter. An ultrafilter of this form is called a *principal ultrafilter*. Zorn's lemma implies that non-principal ultrafilters always exist (provided I is infinite).

Now fix a non-principal ultrafilter $\mathcal{F}$. If $P(i)$ is a statement depending on $i \in I$, say that $P(i)$ holds *almost always* if $\{i \mid P(i)\} \in \mathcal{F}$. For example, if (X, ρ) is a metric space and (x_i) is a sequence indexed by I, write $x_i \to x \in X$ to mean that for any $\epsilon > 0$, $|x_i - x| \le \epsilon$ almost always. With

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

this definition, one can readily check that any bounded sequence in $\mathbb{R}$ has a limit.

Let $(X_i, \rho_i)_i$ be a sequence of metric spaces indexed by our countable set I. Let $\mathbf{X} = \prod_i X_i$; then a point $\mathbf{x} \in \mathbf{X}$ is a sequence (x_i). Fixing a point $\mathbf{a} \in \mathbf{X}$, let $\mathbf{X}^0 = \{\mathbf{x} \in \mathbf{X} \mid \rho_i(a_i, x_i) \text{ is bounded almost always}\}$. This is independent of the choice of $\mathbf{a}$. Given $\mathbf{x}$ and $\mathbf{y}$ in $\mathbf{X}^0$, let $\rho^\infty(\mathbf{x}, \mathbf{y}) = \lim \rho_i(x_i, y_i)$, noting that $\rho_i(x_i, y_i)$ is almost always bounded by the triangle inequality. Then ρ^∞ is a pseudometric on $\mathbf{X}^0$.

Write $\mathbf{x} \sim \mathbf{y}$ if $\rho^\infty(\mathbf{x}, \mathbf{y}) = 0$. Let $X^\infty = \mathbf{X}^0/\sim$. Then ρ^∞ descends to a metric on X^∞. It is then a general fact that X^∞ is complete; this requires that I be countable but not that the X_i be complete.

We are interested in a special case of this definition. Let (X, ρ) be a metric space and let $r_i \geq 0$ be a sequence tending to infinity. Define a new metric ρ_i on X by setting $\rho_i = \rho/r_i$. Then the limit X^∞ of the sequence of spaces (X, ρ_i) is called an *asymptotic cone* of (X, ρ). In general this might depend on the choice of r_i, or the choice of ultrafilter. However, the choice will not matter to us here.

If X is a Gromov hyperbolic space then any asymptotic cone is an $\mathbb{R}$-tree. (This is again a consequence of its treelike structure.) Conversely, a geodesic metric space all of whose asymptotic cones are $\mathbb{R}$-trees is hyperbolic [G2, Dr].

Note that some types of maps between spaces induce maps between their asymptotic cones: if $\phi\colon X \to Y$ is a coarsely-Lipschitz map (respectively a quasi-isometric embedding), then it induces a map $\phi^\infty\colon X^\infty \to Y^\infty$ that is Lipschitz (respectively bi-Lipschitz to its range). This implies, for example, that if there exists a quasi-isometric embedding $\mathbb{R}^n \to X$, then there is a bi-Lipschitz embedding $\mathbb{R}^n \to X^\infty$, so X^∞ has locally compact dimension at least n. (Recall that this is the maximal dimension of any locally compact subset.)

Note that axiom (C1) of a coarse median space Λ tells us that the median operation, μ, is coarsely-Lipschitz and so gives rise to a Lipschitz operation, $\mu^\infty : (\Lambda^\infty)^3 \to \Lambda^\infty$, on its asymptotic cone. In fact, we have the following.

Theorem 1.6.2 *If (Λ, ρ, μ) is a coarse median space, then $(\Lambda^\infty, \rho^\infty, \mu^\infty)$ is a locally convex topological median algebra with* $\text{M-rk}(\Lambda^\infty) \leq \text{C-rk}(\Lambda)$.

Note that Theorem 1.3.7 tells us that Λ^∞ has locally compact dimension at most $\text{M-rk}(\Lambda^\infty)$.

From this we can deduce Theorem 1.4.4, since any quasi-isometric

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

embedding of H^n into Λ would give rise to a continuous (bi-Lipschitz) embedding of H^n into Λ^∞, and so $n \le$ M-rk$(\Lambda^\infty) \le$ C-rk(Λ).

Some arguments will be made simpler if we can assume the metric to be a median metric. The following theorem will allow us to do this.

Theorem 1.6.3 *[Bo3, Bo6] If* C-rk $\Lambda < \infty$*, then* Λ^∞ *is bi-Lipschitz equivalent to a median metric via a median isomorphism.*

(In fact, under slightly stronger hypotheses applicable in the cases of interest to us, one can show that Λ^∞ embeds into a finite direct product of $\mathbb{R}$-trees by a bi-Lipschitz median homomorphism [Bo3].)

Note that by Theorem 1.3.8 we see that Λ^∞ is also bi-Lipschitz equivalent to a CAT(0) metric, and so in particular is contractible.

In general, asymptotic cones have a very complicated structure. However we have the following regularity theorem for median metric spaces. It is based on an analogous result of Kleiner and Leeb [KL].

Theorem 1.6.4 *[Bo6] Let M be a complete median metric space with* M-rk$(M) = n < \infty$. *Suppose that $f: \mathbb{R}^n \to M$ is a continuous injective map with closed image, where n is the rank of M. Then $f(\mathbb{R}^n)$ is cubulated.*

This means that $f(\mathbb{R}^n)$ is a locally finite union of n-dimensional ℓ^1-cubes: each is a convex subset of M isometric (and hence median isomorphic) to an ℓ^1 direct product of n real intervals. In other words, $f(\mathbb{R}^n)$ has the local structure of a cube complex. The complex might still bend along codimension-1 faces. However, this cannot happen if there are lots of other transverse subsets of this form.

Theorem 1.6.5 *[Bo6]. Suppose that M, f are as in Theorem 1.6.4, and suppose, in addition, that for any codimension-1 co-ordinate subspace $P \subset \mathbb{R}^n$, there is another proper embedding $f' : \mathbb{R}^n \to M$ such that $f(P) = f(\mathbb{R}^n) \cap f'(\mathbb{R}^n)$. Then $f(\mathbb{R}^n)$ is convex in M and f is a median homomorphism.*

Here a "codimension-1 co-ordinate subspace" of $\mathbb{R}^n$ is a subset of the form $\{(x_1, \dots, x_n) \mid x_i = t\}$ for some $i \in \{1, \dots, n\}$ and $t \in \mathbb{R}$.

Now let Σ be a compact surface with $\xi(\Sigma) \ge 2$. We consider the case where $\Lambda = \mathcal{M} = \mathcal{M}(\Sigma)$. By Theorem 1.6.3, $\mathcal{M}^\infty$ is bi-Lipschitz equivalent to a median metric, and so Theorems 1.6.4 and 1.6.5 apply with $n = \xi$.

Suppose that α is a pants decomposition, i.e., a collection of ξ disjoint curves on a surface of complexity ξ. Let $T(\alpha) = \{a \in \mathcal{M} \mid \alpha \subset a\}$ be the "Dehn twist flat". (Note that it is a bounded Hausdorff distance from a $\mathbb{Z}^\xi$-orbit, where $\mathbb{Z}^\xi$ is the subgroup generated by Dehn twists about

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

the component curves.) Then the inclusion of $T(\alpha)$ into $\mathcal{M}$ induces an inclusion $T(\alpha)^\infty \subset \mathcal{M}^\infty$. In fact, we get a map $f : \mathbb{R}^\xi \to \mathcal{M}^\infty$ as in Theorem 1.6.4 with $f(\mathbb{R}^\xi) = T(\alpha)^\infty$. It turns out that it also satisfies the hypotheses of Theorem 1.6.5. The basic idea behind this is that one could replace any element $\gamma \in \alpha$ by a different curve γ' so as to give a new pants decomposition, α'. Now $T(\alpha)$ and $T(\alpha')$ remain close near a $\mathbb{Z}^{\xi-1}$-orbit (generated by Dehn twists about the components of $\alpha - \gamma = \alpha' - \gamma'$), and they diverge elsewhere. It then follows that $T(\alpha)^\infty$ meets $T(\alpha')^\infty$ in the f-image of a co-ordinate plane. Elaborating on this idea, one can verify the hypotheses of Theorem 1.6.5.

In fact, we have a converse, which we state informally as follows.

Let $f : \mathbb{R}^\xi \to \mathcal{M}^\infty$ be as in Theorem 1.6.5 (with $n = \xi$ and $M = \mathcal{M}^\infty$).

Theorem 1.6.6 *[Bo6] Sets of the form $f(\mathbb{R}^\xi) \subseteq \mathcal{M}^\infty$ (i.e., as in the hypotheses if Theorem 1.6.5) are precisely the asymptotic Dehn twist flats.*

An example of an "asymptotic Dehn twist flat" a set of the form $T(\alpha)^\infty$ as described above. However, we also need to allow sets constructed by taking I-sequences of pants decompositions rather than just a fixed pants decomposition. The key point here is that we can recognise such sets just in terms of the topology of $\mathcal{M}^\infty$.

Suppose now that $\phi : \mathcal{M} \to \mathcal{M}$ is a quasi-isometry. This induces a (bi-Lipschitz) homeomorphism $\phi^\infty : \mathcal{M}^\infty \to \mathcal{M}^\infty$. By Theorem 1.6.6, we see that ϕ^∞ preserves the collection of asymptotic Dehn twist flats. From this one can go back and deduce that ϕ sends any Dehn twist flat to within a bounded Hausdorff distance of another Dehn twist flat. Now the coarse arrangement of Dehn twist flats in $\mathcal{M}$ can be encoded in terms of the curve graph $\mathcal{C}$. It follows that ϕ gives rise to an automorphism of $\mathcal{C}$. By the combinatorial rigidity result of [I, L, K], this is induced by an element of $\mathrm{Map}(\Sigma)$, which, without loss of generality, we can take to be the identity. We now know that ϕ moves each Dehn twist flat a bounded Hausdorff distance. There are plenty of Dehn twist flats, and it follows easily that ϕ moves each point of $\mathcal{M}$ a bounded distance. This proves the quasi-isometric rigidity of $\mathcal{M}$, as formulated in Theorem 1.5.6.

In fact, we only really need that ϕ is a quasi-isometric embedding. Then ϕ^∞ maps $\mathcal{M}^\infty$ injectively onto a closed subset, which is enough to see that every asymptotic Dehn twist flat gets sent to another such. Following the argument through, this time we get an injection of $\mathcal{C}$ to itself, and the result of [Sha] tells us that it must be an isomorphism,

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press

again induced by $\mathrm{Map}(\Sigma)$. We deduce that ϕ is a quasi-isometry, and close to an element of $\mathrm{Map}(\Sigma)$. This then proves Theorem 1.5.7.

While the details are (significantly) different, related arguments can be made to work for quasi-isometries of $\mathcal{T}$ and $\mathcal{W}$, giving the rigidity results for these spaces [Bo5, Bo7]. We remark that the rigidity of $\mathcal{T}$ is independently proven in [EMR2] using quite different arguments of coarse differentiation.

References

[BaH] H.-J. Bandelt, J. Hedlikova, *Median algebras* : Discrete Math. **45** (1983) 1–30.

[BHS1] J.A. Behrstock, M.F. Hagen, A. Sisto, *Hierarchically hyperbolic spaces I: curve complexes for cubical groups* : Geom. Topol. **21** (2017) 1731–1804.

[BHS2] J.A. Behrstock, M.F. Hagen, A. Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula.* : to appear in Pacific J. Math.

[BKMM] J.A. Behrstock, B. Kleiner, Y. Minsky, L. Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.

[BM1] J.A. Behrstock, Y.N. Minsky, *Dimension and rank for mapping class groups* : Ann. of Math. **167** (2008) 1055–1077.

[BM2] J.A. Behrstock, Y.N. Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.

[Bo1] B.H. Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.

[Bo2] B.H. Bowditch, *Invariance of coarse median spaces under relative hyperbolicity* : Math. Proc. Camb. Phil. Soc. **154** (2013) 85–95.

[Bo3] B.H. Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.

[Bo4] B.H. Bowditch, *Some properties of median metric spaces* : Groups, Geom. Dyn. **10** (2016) 279–317.

[Bo5] B.H. Bowditch, *Large-scale rank and rigidity of the Teichmüller metric* : J. Topol. **9** (2016) 985–1020.

[Bo6] B.H. Bowditch, *Large-scale rigidity properties of the mapping class groups* : Pacific J. Math. **293** (2018), no. 1, 1–73.

[Bo7] B.H. Bowditch, *Large-scale rank and rigidity of the Weil–Petersson metric* : preprint, Warwick, 2015.

[Bo8] B.H. Bowditch, *Median and injective metric spaces* : to appear in Math. Proc. Cam. Phil. Soc.

[BrH] M. Bridson, A. Haefliger, *Metric spaces of non-positive curvature* : Grundlehren der Math. Wiss. No. 319, Springer (1999).

[Br] J.F. Brock, *The Weil–Petersson metric and volumes of 3-dimensional hyperbolic convex cores* : J. Amer. Math. Soc. **16** (2003) 495–535.

[Dr] C. Druţu, *Quasi-isometry invariants and asymptotic cones* :in "International Conference on Geometric and Combinatorial Methods in Group Theory and Semigroup Theory" : Internat. J. Algebra Comput. **12** (2002) 99–135.

[Du1] M.G. Durham, *The augmented marking complex of a surface* : J. London Math. Soc. **94** (2016) 933–969.

[Du2] M.G. Durham, *The asymptotic geometry of the Teichmüller metric: Dimension and rank* : preprint, 2014, posted at arXiv:1501.00200.

[E] R. Engelking, *Dimension theory* : North Holland (1978).

[EMR1] A. Eskin, H. Masur, K. Rafi, *Large scale rank of Teichmüller space* : Duke. Math. J. **166** (2017) 1517–1572.

[EMR2] A. Eskin, H. Masur, K. Rafi, *Rigidity of Teichmüller space* : Geom. Topol. **22** (2018) 4259–4306.

[G1] M. Gromov, *Hyperbolic groups* : in "Essays in group theory" Math. Sci. Res. Inst. Publ. No. 8, Springer (1987) 75–263.

[G2] M. Gromov, *Asymptotic invariants of infinite groups* : "Geometric group theory, Vol. 2" London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993).

[H] U. Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity* : preprint, 2007, posted at arXiv:0512429.

[IT] Y. Imayoshi, M. Taniguchi, *An introduction to Teichmüller spaces* : Springer Verlag (1992).

[I] N.V. Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces* : Internat. Math. Res. Notices **14** (1997) 651–666.

[KR] M. Kapovich, K. Rafi, *Teichmüller space is semi-hyperbolic* : in preparation.

[KL] B. Kleiner, B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings* : Inst. Hautes Études Sci. Publ. Math. **86** (1997) 115–197.

[K] M. Korkmaz, *Automorphisms of complexes of curves in punctured spheres and on punctured tori* : Topology Appl. **95** (1999) 85–111.

[L] F. Luo, *Automorphisms of the complex of curves* : Topology **39** (2000) 283–298.

[MM1] H.A. Masur, Y.N. Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103–149.

[MM2] H.A. Masur, Y.N. Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.

[Mi] B. Miesch, *Injective metric on median metric spaces* : in preparation.

[Mo] L. Mosher, *Mapping class groups are automatic* : Ann. Math. **142** (1995) 303–384.

[NWZ] G.A. Niblo, N. Wright, J. Zhang, *A four point characterisation for coarse median spaces* : to appear in Groups. Geom. Dyn.

[Ra] K. Rafi, *A combinatorial model for the Teichmüller metric* : Geom. Funct. Anal. **17** (2007) 936–959.

[RS] K. Rafi, S. Schleimer, *Curve complexes are rigid* : Duke Math J. **158** (2011) 225–246.

[Ro] M.A. Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody's construction and Sageev's theorem* : Habilitationschrift, Regensberg, 1998.

[Sha] K.J. Shackleton, *Combinatorial rigidity in curve complexes and mapping class groups* : Pacific J. Math. **230** (2007) 217–232.

[Sho] M. Sholander, *Medians and betweenness* : Proc. Amer. Math. Soc. **5** (1954) 801–807.

[vdDW] L. van den Dries, A.J. Wilkie, *On Gromov's theorem concerning groups of polynomial growth and elementary logic* : J. Algebra **89** (1984) 349–374.

[Ve] E.R. Verheul, *Multimedians in metric and normed spaces* : CWI Tract, 91, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam (1993).

[Vo] K.M. Vokes, *Hierarchical hyperbolicity of graphs of multicurves* : posted at arXiv:1711.03080.

[W] S. Wolpert, *Noncompleteness of the Weil–Petersson metric for Teichmüller space* : Pacific J. Math. **61** (1975) 573–577.

https://doi.org/10.1017/9781108559065.003 Published online by Cambridge University Press