Cambridge University Press
978-1-108-44729-4 — Beyond Hyperbolicity
Edited by Mark Hagen , Richard Webb , Henry Wilton
Excerpt
More Information

PART ONE

LECTURES

Cambridge University Press
978-1-108-44729-4 — Beyond Hyperbolicity
Edited by Mark Hagen , Richard Webb , Henry Wilton
Excerpt
More Information

Cambridge University Press
978-1-108-44729-4 — Beyond Hyperbolicity
Edited by Mark Hagen , Richard Webb , Henry Wilton
Excerpt
More Information

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.

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 $X \not\sim Y$. 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

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

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.)*

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 $\phi\colon I \to X$ 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 $\mathbb{R}^2 \not\leq \mathbb{R}$.

Proof Suppose for contradiction that $\phi\colon \mathbb{R}^2 \to \mathbb{R}$ 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 $H^2 \not\leq \mathbb{R}$. Moreover, replacing the Intermediate Value Theorem with the Borsuk–Ulam theorem, one can see that $\mathbb{R}^{n+1} \not\leq \mathbb{R}^n$ for any n, and therefore $\mathbb{R}^m \sim \mathbb{R}^n$ only when $m = n$. Indeed one can see that $H^{n+1} \not\leq \mathbb{R}^n$. 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 $\text{E-rk}(X) \in \mathbb{N} \cup \{\infty\}$ is defined to be the maximum n such that $\mathbb{R}^n \leq X$. The *half-space rank* of X, $\text{H-rk}(X)$, is defined to be the maximum n such that $H^n \leq X$.

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

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

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.)

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.

Cambridge University Press
978-1-108-44729-4 — Beyond Hyperbolicity
Edited by Mark Hagen , Richard Webb , Henry Wilton
Excerpt
More Information

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 $B \le M$. For any $A \subset M$, $\langle A \rangle \le M$ 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 $|A| \le p < \infty$. Then* $|\langle A \rangle| \le 2^{2^p}$.

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 $\{0,1\}^n \le M$, so M-rk$(M) \in \mathbb{N} \cup \{\infty\}$.