EMBEDDINGS OF GROMOV HYPERBOLIC SPACES

M. BONK AND O. SCHRAMM

Abstract

It is shown that a Gromov hyperbolic geodesic metric space X with bounded growth at some scale is roughly quasi-isometric to a convex subset of hyperbolic space. If one is allowed to rescale the metric of X by some positive constant, then there is an embedding where distances are distorted by at most an additive constant.

Another embedding theorem states that any δ-hyperbolic metric space embeds isometrically into a complete geodesic δ-hyperbolic space.

The relation of a Gromov hyperbolic space to its boundary is further investigated. One of the applications is a characterization of the hyperbolic plane up to rough quasi-isometries.

1 Introduction

The study of Gromov hyperbolic spaces has been largely motivated and dominated by questions about Gromov hyperbolic groups. This paper studies the geometry of Gromov hyperbolic spaces without reference to any group or group action.

One of our main theorems is

Embedding Theorem 1.1. *Let X be a Gromov hyperbolic geodesic metric space with bounded growth at some scale. Then there exists an integer n such that X is roughly similar to a convex subset of hyperbolic n-space $\mathbb{H}^n$.*

The precise definitions appear in the body of the paper, but now we briefly discuss the meaning of the theorem. The condition of bounded growth at some scale is satisfied, for example, if X is a bounded valence graph, or a Riemannian manifold with bounded local geometry. Saying

that X is roughly similar to a metric space Y means that there is a map $f : X \to Y$ and constants $k, \lambda > 0$, such that

$$\big| \lambda d_X(x, y) - d_Y(f(x), f(y)) \big| \leqslant k, \qquad \forall x, y \in X,$$

and $\sup_{y \in Y} d_Y(y, f(X)) \leqslant k$. This is much stronger than saying that X and Y are roughly quasi-isometric. (Roughly quasi-isometric is the term we use for what is elsewhere "quasi-isometric" or "roughly-isometric".)

It follows in particular that complex hyperbolic m-space embeds in $\mathbb{H}^n$ for some $n = n(m)$. It would be interesting to determine or estimate the smallest possible $n(m)$.

The converse of the theorem is straightforward: a geodesic metric space X that is roughly similar to a convex subset W of $\mathbb{H}^n$ has bounded growth at some scale and is Gromov hyperbolic. This is even known to be true if you weaken the assumption that X is roughly similar to W and instead assume that X is roughly quasi-isometric to W.

The outline of the proof of Theorem 1.1 is quite simple. We verify that the Gromov boundary ∂X has finite Assouad dimension, and then Assouad's theorem [A] implies that (if we rescale the metric of X by a constant) ∂X has a bilipschitz embedding in $\mathbb{R}^{n-1}$, for some n. As $\mathbb{R}^{n-1} \cup \{\infty\}$ is the boundary of $\mathbb{H}^n$, further assuming that the set of all geodesics is cobounded in X (that is, there is a finite upper bound for the distance of any point in X to the union of the geodesics in X), the embedding of ∂X in $\partial \mathbb{H}^n$ induces an embedding of X in $\mathbb{H}^n$. The general statement then easily reduces to this case.

Another embedding theorem presented here says that any δ-hyperbolic metric space X embeds isometrically in a complete geodesic δ-hyperbolic metric space. This allows us to generalize statements about geodesic Gromov hyperbolic metric spaces to the non-geodesic setting.

It is known that a rough quasi-isometry $f : X \to Y$ between Gromov hyperbolic metric spaces induces a map $\partial f : \partial X \to \partial Y$, which is bi-hölder and quasiconformal. However, it turns out that ∂f satisfies a stronger condition. It is what is called below a power quasisymmetry. Conversely, under mild conditions on X and Y, any power quasisymmetry from ∂X to ∂Y induces a rough quasi-isometry from X to Y.

Previously, F. Paulin [P] has obtained a characterization of X from ∂X under the assumption that there is a group acting isometrically and co-compactly on X. His methods and results are different.

Given any bounded metric space Z, we define a Gromov hyperbolic space $\mathrm{Con}(Z)$ such that $\partial \mathrm{Con}(Z) = Z$. This is quite similar to a construc-

tion by Gromov. It turns out that under mild assumptions $\mathrm{Con}(\partial X)$ is roughly quasi-isometric to X, when X is Gromov hyperbolic.

As an application of the machinery developed, some results about two dimensional Gromov hyperbolic spaces are obtained. We give characterization of $\mathbb{H}^2$ up to rough quasi-isometries.

It is also shown that a planar Gromov hyperbolic graph with bounded vertex degree and finitely many vertices in any compact subset of the plane is roughly quasi-isometric to a convex subset of $\mathbb{H}^2$, provided that there is a finite upper bound for the distance of a vertex in G to the union of all geodesics in G.

The paper is organized as follows. After reviewing some notions of coarse geometry in section 2, and some basics of Gromov hyperbolic spaces in section 3, we show in section 4 that each Gromov hyperbolic space can be embedded in a complete geodesic Gromov hyperbolic space without changing the hyperbolicity constant δ. In section 5 we collect some auxiliary results about Gromov hyperbolic spaces that are needed later on.

In sections 6–8 we discuss the boundary functor $X \mapsto \partial X$, the convex hull functor $Z \mapsto \mathrm{Con}(Z)$, and their relation. This is the technical core of the paper. We give particular emphasis to the correspondence of maps between Gromov hyperbolic spaces and maps between their boundaries.

In section 9 we recall the notion of Assouad dimension, and establish a sufficient condition for a Gromov hyperbolic space to have a boundary with finite Assouad dimension.

The last two sections give applications of the methods developed in earlier sections. In section 10 Theorem 1.1 is proved, and in section 11 the real hyperbolic plane is characterized up to rough quasi-isometry among geodesic metric spaces.

He is pleased to acknowledge the hospitality of the local Department of Mathematics.

2   Coarse Geometry Definitions

Rough quasi-isometries.   Let (X, d) be a metric space. If the metric d is fixed, we denote the distance $d(x, y)$ of two points $x, y \in X$ also by $|x - y|$. The diameter of a set $A \subset X$ is denoted by $\mathrm{diam}(A)$, and the distance of two set $A, B \subset X$ by $\mathrm{dist}(A, B)$.

A set $A \subset X$ is *k-cobounded* (in X) for $k \geqslant 0$, if every point $x \in X$ has distance at most k from A. If A is k-cobounded for some $k \geqslant 0$ we say that A is *cobounded*. In the same way, we will drop the parameters of a notion if the values of the parameters do not matter.

Let $f : X \to Y$ be a map (not necessarily continuous) between metric spaces X and Y, and let $\lambda \geqslant 1$ and $k \geqslant 0$ be constants. Suppose that $f(X)$ is k-cobounded in Y. If, in addition, for all $x, y \in X$,

$$\lambda^{-1}|x - y| - k \leqslant \big|f(x) - f(y)\big| \leqslant \lambda|x - y| + k\,, \tag{2.1}$$

then f is called a (λ, k)-*rough quasi-isometry*. If

$$\lambda|x - y| - k \leqslant \big|f(x) - f(y)\big| \leqslant \lambda|x - y| + k\,, \tag{2.2}$$

then f is a (λ, k)-*rough similarity*. If

$$|x - y| - k \leqslant \big|f(x) - f(y)\big| \leqslant |x - y| + k\,, \tag{2.3}$$

then f is a *k-rough isometry*.

If $f : X \to Y$ is a map that satisfies (2.1) for all $x, y \in X$ without $f(X)$ being cobounded in Y, then we call f a (λ, k)-rough quasi-isometry of X *into* Y or a *rough quasi-isometric embedding*. We use similar language for a map that satisfies (2.2) or (2.3), but does not necessarily have cobounded image in its target space. Each rough isometry is also a rough similarity, and each rough similarity is a rough quasi-isometry. If f is bijective and satisfies (2.1) for $k = 0$, then it is called a λ-*quasi-isometry* or λ-*bilipschitz* map.

The classes of maps just introduced have appeared under various names in the literature. Here and in the following we have tried to adopt the following system for our terminology. The word "rough" refers to change of an equality, inequality, etc., defining a notion by an additive constant, and the word "quasi" to a change by a multiplicative constant. Keeping this in mind, the above terms are more or less self-explanatory. We sometimes deviate from this system, if we feel that a certain term is well-established.

Two maps $f, g : X \to Y$ are *roughly equivalent*, written $f \approx g$, if there exists some constant $k \geqslant 0$ such that $|f(x) - g(x)| \leqslant k$ for $x \in X$. This defines an equivalence relation on the set of mappings from X to Y. The equivalence class of a map f is called the *rough mapping class* of f. If $f_1, f_2 : X \to Y$ and $g_1, g_2 : Y \to Z$ are rough quasi-isometric embeddings, then $f_1 \approx f_2$ and $g_1 \approx g_2$ imply $g_1 \circ f_1 \approx g_2 \circ f_2$. Compositions of rough quasi-isometric embeddings are maps of the same type. The same statement is true for embeddings by rough similarities or rough isometries. From this it follows that it is possible to define categories, where the objects are metric spaces and the morphisms are rough mapping classes of our various types of embeddings.

A *rough-inverse* of a rough quasi-isometry $f : X \to Y$ is a rough quasi-isometry $g : Y \to X$ such that $g \circ f \approx \mathrm{id}_X$ and and $f \circ g \approx \mathrm{id}_Y$. Every rough quasi-isometry has a rough inverse. If in addition the rough quasi-isometry is a rough similarity or a rough isometry, then the rough inverse is also a rough similarity or a rough isometry, respectively.

We say that two metric spaces X and Y are *roughly quasi-isometric*, written $X \sim Y$, if there exists a rough quasi-isometry $f : X \to Y$. This is an equivalence for the class of metric spaces. In a similar way we define that X and Y are *roughly similar*, $X \simeq Y$, or *roughly isometric*, $X \cong Y$. These are also equivalences. The equivalence $X \cong Y$ implies $X \simeq Y$, and $X \simeq Y$ implies $X \sim Y$.

Geodesics and rough geodesics in metric spaces.   Let X be a metric space. A *path* in X is a map $\gamma : I \to X$, where I is an interval in $\mathbb{R}$. We will also denote by γ the image set of the path γ. A path γ *joins* two points x and y in X, if $I = [a, b]$, $\gamma(a) = x$ and $\gamma(b) = y$. The path γ is called a *ray starting*, or *emanating*, from the *initial point* x, if $I = [0, \infty)$ and $\gamma(0) = x$. If $\gamma : I \to X$ is a (λ, k)-rough quasi-isometry of I into X, then γ is called a (λ, k)-*roughly quasi-isometric path*.

A *geodesic*, a *geodesic ray* or a *geodesic segment* in X is an isometry $\gamma : I \to X$ into X, where I is $\mathbb{R}$, $[0, \infty)$ or a closed segment in $\mathbb{R}$, respectively. Similarly, if $\gamma : I \to X$ is a k-rough isometry of $I = \mathbb{R}$ into X, we call it a k-*rough geodesic*. We also speak of k-*roughly geodesic rays* or k-*roughly geodesic segments*, if $I = [0, \infty)$ or I is a closed segment in $\mathbb{R}$, respectively.

We call subsets of X geodesics, rough geodesics, etc., if they are images of a path in the corresponding class.

A *geodesic metric space* is a metric space X such that for any two points $x, y \in X$ there is a geodesic segment joining x and y. We denote

any geodesic segment with endpoints x, y by $[x, y]$. Note that this notation is ambiguous, since $[x, y]$ may not be unique.

The space X is *k-roughly geodesic*, if for every $x, y \in X$ there exists a k-roughly geodesic segment joining x and y. It is *k-almost geodesic*, if for every $x, y \in X$ and every $t \in [0, |x - y|]$, there is some $z \in X$ with $||x - z| - t| \leqslant k$ and $||y - z| - (|x - y| - t)| \leqslant k$.

The reader, who might be content with studying only geodesic metric spaces, may question the motivation for introducing the concepts of almost geodesic and roughly geodesic. It turns out that the setting of roughly geodesic spaces is more fitting with the philosophy of Gromov hyperbolic spaces, and it is worthwhile to work in that setting. In particular, the spaces $\mathrm{Con}(Z)$, which we will construct and will be important for us, are roughly geodesic, and it will be unnecessarily cumbersome to modify the construction to make them truly geodesic. Remark 5.3 below further discusses this issue.

Constants. In the following, we will encounter various constants whose precise values usually does not matter.

The letter C will denote positive numerical constants. Similarly, $C(a, b, c, \ldots)$ will denote universal positive functions of the parameters $a, b, c, \ldots$. Sometimes we write $C = C(a, b, c, \ldots)$ to emphasize the parameters on which $C(a, b, c, \ldots)$ depends and abbreviate $C(a, b, c, \ldots)$ to C. We will use the same letter for possibly different constants and functions.

The notation $O(1)$ will be used for real numbers whose absolute value is bounded by a positive numerical constant that can be computed explicitly in principle. Similarly, we use the notation $O_{a,b,c,\ldots}(1)$ for real numbers whose absolute value is bounded by a constant which can be chosen in a way only depending on the parameters $a, b, c, \ldots$.

3   Some Basics of Gromov Hyperbolic Spaces

We here review some definitions and elementary facts concerning Gromov hyperbolic spaces. For a more detailed account of this material, the reader is referred to [CDP], [GH] or [Gr].

The definition of Gromov hyperbolicity. Let X be a metric space. Given three points $x, y, w \in X$, the *Gromov product* of x and y with respect to the basepoint w is defined as

$$(x|y)_w = \tfrac{1}{2}\left(|x - w| + |y - w| - |x - y|\right).$$

It measures the failure of the triangle inequality to be an equality, and is always nonnegative. Often we will fix a basepoint $w = o \in X$. Then we abbreviate the Gromov product as $(x|y) = (x|y)_o$, and let $|x| = |x - o|$ denote the distance of x from the basepoint. For all $x, y, v, w \in X$ we have

$$\left| (x|y)_v - (x|y)_w \right| \leqslant |v - w| . \tag{3.1}$$

This formula is tacitly used in the following when it is convenient to change the basepoint.

Let $a \vee b$ denote the maximum of two numbers $a, b \in \overline{\mathbb{R}}$, and let $a \wedge b$ denote the minimum. The space X is δ-*hyperbolic*, if the following inequality holds for all $x, y, z, w \in X$,

$$(x|z)_w \geqslant (x|y)_w \wedge (y|z)_w - \delta . \tag{3.2}$$

This is equivalent to the more symmetric inequality

$$|x - z| + |y - w| \leqslant \left(|x - y| + |z - w| \right) \vee \left(|y - z| + |x - w| \right) + 2\delta . \tag{3.3}$$

It turns out that if (3.2) holds for fixed $w \in X$ and arbitrary $x, y, z \in X$, then X is 2δ-hyperbolic. If X is δ-hyperbolic for some $\delta < \infty$, we say that X is *Gromov hyperbolic*.

The Rips thin triangles definition of Gromov hyperbolicity. Let X be a geodesic metric space. We say that X has δ-*thin triangles*, if for every geodesic triangle $[x, y] \cup [y, z] \cup [z, x] \subset X$ and every point $w \in [x, y]$, the distance from w to $[y, z] \cup [z, x]$ is at most δ. There exists some constant $C > 0$ with the following property. If X has δ-thin triangles, then X is $C\delta$-hyperbolic, and if X is δ-hyperbolic, then it has $C\delta$-thin triangles. Hence, in the context of geodesic metric spaces, the condition of having thin triangles provides an alternative definition for Gromov hyperbolicity. This definition was introduced by I. Rips.

The boundary of a Gromov hyperbolic space. Let X be δ-hyperbolic and $w \in X$. A sequence of points $\{x_i\} \subset X$ is said to *converge at infinity*, if

$$\lim_{i,j \to \infty} (x_i | x_j)_w = \infty .$$

Two sequences $\{x_i\}, \{y_i\}$ that converge at infinity are *equivalent*, if

$$\lim_{i \to \infty} (x_i | y_i)_w = \infty .$$

This defines an equivalence relation for sequences in X converging at infinity. It is easy to see that convergence at infinity of a sequence and equivalence of two sequences does not depend on the choice of the point w. The *boundary* ∂X of X is defined as the set of equivalence classes of sequences converging at infinity.

Let $a, b \in \partial X$. The Gromov product $(a|b)_w$ of a and b with respect to the point $w \in X$ is defined as follows:

$$(a|b)_w = \sup \left\{ \liminf_{i \to \infty} (x_i|y_i)_w \;\middle|\; \{x_i\} \in a \,, \; \{y_i\} \in b \right\} .$$

Similarly, for $a \in \partial X$ and $y \in X$ we define

$$(a|y)_w = (y|a)_w = \sup \left\{ \liminf_{i \to \infty} (x_i|y)_w \;\middle|\; \{x_i\} \in a \right\} .$$

Note that $(a|b)_w \in [0, \infty]$ for $a, b \in X \cup \partial X$, and $(a|b)_w = \infty$ if and only if $a, b \in \partial X$, $a = b$. It can be shown that

$$(a|b)_w - C\delta \leqslant \liminf_{i \to \infty} (x_i|y_i)_w \leqslant \limsup_{i \to \infty} (x_i|y_i)_w \leqslant (a|b)_w + C\delta \,,$$

$$\forall \, a, b \in \partial X \,, \; \{x_i\} \in a \,, \; \{y_i\} \in b \,. \quad (3.4)$$

A similar inequality holds if one of the points a, b is in X. Inequality (3.2) generalizes to points in ∂X; more precisely,

$$(x|z)_w \geqslant (x|y)_w \wedge (y|z)_w - C\delta \,, \quad \forall \, w \in X \,, \; x, y, z \in X \cup \partial X \,. \quad (3.5)$$

Geodesic stability. If $A \subset X$ and $B \subset X$ are two subsets of a metric space X, define their *Hausdorff distance* to be

$$d_{\mathrm{Haus}}(A, B) = \left(\sup_{a \in A} \left(\inf_{b \in B} |a - b| \right) \right) \vee \left(\sup_{b \in B} \left(\inf_{a \in A} |a - b| \right) \right) .$$

Suppose X is a δ-hyperbolic geodesic metric space, and let $\lambda \geqslant 1$, $k \geqslant 0$. Then there exists a constant $C = C(\delta, \lambda, k)$ with the following property. If $\gamma : [a, b] \to X$ is a (λ, k)-roughly quasi-isometric path, $x = \gamma(a)$ and $y = \gamma(b)$ are the endpoints of γ, and $[x, y]$ is any geodesic segment joining x and y, then $d_{\mathrm{Haus}}(\gamma, [x, y]) \leqslant C$.

4   Embedding into a Geodesic Metric Space

Theorem 4.1. *Let X be a δ-hyperbolic metric space. Then there is an isometric embedding $i : X \to Y$ of X into a complete δ-hyperbolic geodesic metric space Y.*

The following lemma provides the basic step in the construction of Y and i.

LEMMA 4.2.   *Let X be a δ-hyperbolic metric space, let $a, b \in X$, and set $D = |a - b|$. Then there is an isometric embedding i of X into a δ-hyperbolic metric space $(\widehat{X}, d)$ which contains a point m satisfying $d(m, i(a)) = d(m, i(b)) = D/2$.*

The main point is that the hyperbolicity constant δ of $\widehat{X}$ is the same as that of X.

Proof. Assume that there is no point $x \in X$ satisfying $|x - a| = |x - b| = D/2$. Let $\widehat{X} = X \cup \{m\}$, where m is a new point not appearing in X. When $x, y \in X$, define $d(x, y) = |x - y|$, and for any $x \in X$, let

$$d(x, m) = d(m, x) = \tfrac{D}{2} + \sup_{w \in X} \left(|x - w| - |a - w| \vee |b - w| \right). \quad (4.1)$$

Note that this is always finite. Finally, let $d(m, m) = 0$.

We now verify that d is a metric on $\widehat{X}$. Let $x, y \in X$. When substituting a and b for w in (4.1), we find

$$d(x, m) \geqslant D/2 + \left(|x - a| - |a - a| \vee |b - a| \right) \vee \left(|x - b| - |a - b| \vee |b - b| \right)$$
$$= |x - a| \vee |x - b| - D/2. \quad (4.2)$$

Since $D \leqslant |x - a| + |x - b|$, and there is no point in X satisfying $|x - a| = |x - b| = D/2$, the right hand side of (4.2) must be positive. Consequently, $d(z, z') > 0$ when $z \neq z'$ are in $\widehat{X}$. Now (4.2) is used to verify one case in the triangle inequality for d.

$$d(x, m) + d(m, y) \geqslant |x - a| \vee |x - b| + \sup_{w \in X} \left(|y - w| - |a - w| \vee |b - w| \right)$$
$$\geqslant |x - a| \vee |x - b| + \left(|y - x| - |a - x| \vee |b - x| \right)$$
$$= |x - y| = d(x, y).$$

Another case of the triangle inequality is proved as follows,

$$d(m, x) + d(x, y) = D/2 + \sup_{w \in X} \left(|x - w| - |a - w| \vee |b - w| \right) + |x - y|$$
$$\geqslant D/2 + \sup_{w \in X} \left(|y - w| - |a - w| \vee |b - w| \right)$$
$$= d(m, y).$$

Since the triangle inequality for d holds for any three points in X, it now follows that it holds for any three points in $\widehat{X}$, and hence $\widehat{X}$ is a metric space.

We now verify that $(\widehat{X}, d)$ is δ-hyperbolic. For this, it is enough to prove the inequality (3.3) for m and points $x, y, z \in X$. For an arbitrary $w \in X$, we have

$$|x - w| + |y - z| \leqslant \left(|y - w| + |x - z| \right) \vee \left(|z - w| + |x - y| \right) + 2\delta.$$

Therefore,

$$d(m, x) + d(y, z) = D/2 + \sup_{w \in X} \left(|x - w| - |a - w| \vee |b - w| \right) + |y - z|$$

$$\leqslant D/2 + 2\delta$$
$$+ \sup_{w \in X} \big((|y-w| + |x-z|) \vee (|z-w| + |x-y|) - |a-w| \vee |b-w| \big)$$
$$= \big(D/2 + |x - z| + \sup_{w \in X} (|y - w| - |a - w| \vee |b - w|) \big)$$
$$\vee \big(D/2 + |x - y| + \sup_{w \in X} (|z - w| - |a - w| \vee |b - w|) \big) + 2\delta$$
$$= \big(d(m,y) + d(x,z) \big) \vee \big(d(m,z) + d(x,y) \big) + 2\delta \,,$$

which verifies that $(\widehat{X}, d)$ is δ-hyperbolic.

Observe that
$$d(m,a) = D/2 + \sup_{w \in X} \big(|a - w| - |a - w| \vee |b - w| \big)$$
$$= D/2 + \sup_{w \in X} 0 \wedge \big(|a - w| - |b - w| \big) \,.$$

As b may be chosen for w, this shows that $d(m,a) = D/2$. Similarly, $d(m,b) = D/2$, and the proof is complete. $\qquad\square$

Proof of 4.1. Suppose Z is a metric space and $x, y \in Z$. We say that m is a midpoint of x and y if $|x - m| = |y - m| = |x - y|/2$. The space Z has the midpoint property, if for all $x, y \in Y$ there exists a midpoint $m \in Z$. In the following, when a metric space Z_1 admits an isometric embedding into another metric space Z_2, we can think of Z_1 as being contained in Z_2 and write $Z_1 \subset Z_2$.

The first step in the proof is to construct a metric space $X' \supset X$, which has the midpoint property and is δ-hyperbolic.

For a δ-hyperbolic metric space Z and $a, b \in Z$, denote by $Z[a, b]$ the space constructed according to the previous lemma. Let $\phi : \omega \to X \times X$ be a bijective mapping from some ordinal ω onto the set of pairs in X. By transfinite induction we define a metric space $X(\alpha)$ for each ordinal $\alpha \leqslant \omega + 1$ with the property that $X(\alpha) \supset X(\beta)$ if $\alpha > \beta$.

Set $X(0) = X$. When $\alpha = \beta + 1$ is a successor ordinal, let $X(\alpha)$ be $X(\beta)[a, b]$, where $(a, b) = \phi(\beta)$. For limit ordinals α, set $X(\alpha) = \cup_{\beta < \alpha} X(\beta)$. Note that nested unions of δ-hyperbolic spaces are δ-hyperbolic. It follows that $X^* = X(\omega + 1)$ is a δ-hyperbolic metric space which contains X and has the property that for any two points in X there exists a midpoint in X^*.

Now define inductively $X_0 = X$, and $X_{i+1} = X_i^*$ for $i \in \mathbb{N}$. Let $X' = \cup_{i \in \mathbb{N}} X_i$. Then X' has the required properties.

To construct a complete metric space $Y \supset X$ which has the midpoint property and is δ-hyperbolic we use transfinite induction again. For this,

let ω_0 be the first uncountable ordinal.

We define a metric spaces $Z(\alpha)$ for each ordinal $\alpha < \omega_0$ with the property that $Z(\alpha) \supset Z(\beta)$ if $\alpha > \beta$. Let $Z(0)$ be the completion of X'. When $\alpha = \beta + 1$ is a successor ordinal, let $Z(\alpha)$ be the completion of $Z(\beta)'$. For limit ordinals α, let $Z(\alpha)$ be the completion of $\cup_{\beta<\alpha}Z(\beta)'$. Note that completions of δ-hyperbolic spaces are δ-hyperbolic. It follows by transfinite induction that for $\alpha < \omega_0$ all spaces $Z(\alpha)$ are δ-hyperbolic, complete and have the property that if $x, y \in Z(\alpha)$, then $Z(\alpha + 1)$ contains a midpoint of x and y.

Define $Y = \cup_{\alpha<\omega_0}Z(\alpha)$. By the above properties of the spaces $Z(\alpha)$, this space contains X, is δ-hyperbolic and has the midpoint property. Furthermore, Y is complete. For if $\{x_i\}$ is a Cauchy-sequence in Y, then each element x_i is contained in some $Z(\alpha_i)$ where $\alpha_i < \omega_0$. Since ω_0 is the first uncountable ordinal, there exits $\beta < \omega_0$ such that $\alpha_i \leqslant \beta$ for all $i \in \mathbb{N}$. Hence the Cauchy sequence $\{x_i\}$ is contained in $Z(\beta)$. So it has a limit in the complete space $Z(\beta) \subset Y$.

The space Y is geodesic. This follows from the midpoint property and the completeness of Y. For if $x, y \in Y$ are arbitrary and $D = |x - y|$, then by induction it can be shown that there exists an isometry ϕ from the set $S = \{kD/2^n : n, k \in \mathbb{N}, 0 \leqslant k \leqslant 2^n\}$ into Y with $\phi(0) = x$ and $\phi(D) = y$. By the completeness of Y, ϕ extends to an isometry of $[0, D]$ into Y. The image of this isometry is a geodesic segment $[x, y]$. We leave the details to the reader.

Thus, Y has all the required properties and the proof is complete. $\square$

5 Rough Geodesics in Almost Geodesic Metric Spaces

In this section, we collect various geometric statements that will be used later. Almost all of this material has already appeared in the literature, perhaps in slightly different form. Therefore, we will mostly skip the proofs or only outline the main ideas. An important observation that often serves as a guiding principle in proving statements about a δ-hyperbolic space X is the fact that given any n-point set $M \subset X$ there exists a k-rough isometric embedding of M into a metric tree with $k = k(\delta, n)$ (cf. [GH, pp. 33–38]). In this way geometric statements about finite point sets in X can be reduced to an essentially combinatorial problem about the location of a finite point set in a tree.

We start with a lemma that is very useful to estimate distances of points

in a Gromov hyperbolic space. An intuitive interpretation of this statement is given below.

LEMMA 5.1.   *Let X be a δ-hyperbolic metric space. Let o be the basepoint of X, let $a, b \in X \cup \partial X$, and let $x, y \in X$. Suppose that for some $c \geqslant 0$ we have $(a|x) \geqslant |x| - c$ and $(b|y) \geqslant |y| - c$. Then*

$$|x - y| = |x| + |y| - 2\big((a|b) \wedge |x| \wedge |y|\big) + O_{\delta,c}(1).$$

Proof. On the one hand, (3.5) gives

$$(a|b) \geqslant (a|x) \wedge (x|b) - C\delta$$
$$\geqslant (a|x) \wedge (x|y) \wedge (y|b) - C\delta \geqslant |x| \wedge (x|y) \wedge |y| + O_{\delta,c}(1).$$

Since $(x|y) \leqslant |x| \wedge |y|$, this implies $|x| \wedge (a|b) \wedge |y| \geqslant (x|y) + O_{\delta,c}(1)$.
   On the other hand,

$$(x|y) \geqslant (x|a) \wedge (a|y) - C\delta$$
$$\geqslant (x|a) \wedge (a|b) \wedge (b|y) - C\delta \geqslant |x| \wedge (a|b) \wedge |y| + O_{\delta,c}(1).$$

The combination of these two inequalities is

$$(x|y) = |x| \wedge (a|b) \wedge |y| + O_{\delta,c}(1).$$

The lemma follows, because $|x - y| = |x| + |y| - 2(x|y)$.   □

A roughly geodesic ray $\gamma : [0, \infty) \to X$ in a Gromov hyperbolic space *converges at infinity*, in the following sense. For each sequence $\{t_i\} \to \infty$, the sequence $\{\gamma(t_i)\}$ converges at infinity, and the equivalence class of this sequence does not depend on the choice of $\{t_i\}$. The equivalence class of $\{\gamma(t_i)\}$ is called the *endpoint* of γ on ∂X, or the *limit* of γ. We will denote it by $\lim_{t \to \infty} \gamma(t)$ or simply by $\lim \gamma$. The endpoints of a rough geodesic $\gamma : \mathbb{R} \to X$ are defined similarly.

Suppose γ_1 and γ_2 are two roughly geodesic rays in a Gromov hyperbolic space X with initial point o and endpoints a and b, respectively. Suppose $a \neq b$. We can apply Lemma 5.1 for points $x = \gamma_1(t)$ and $y = \gamma_2(t)$, $t \geqslant 0$. In this situation the lemma says that the roughly geodesic rays run close to each other as long as $0 \leqslant t \leqslant (a|b)$, but they start to spread at $t = (a|b)$.

PROPOSITION 5.2.   *Let X be a δ-hyperbolic, k-almost geodesic metric space. Then there exists a constant $k' = k'(\delta, k)$ with the following properties:*

(1) *For all points $x, y \in X$, there exists a k'-roughly geodesic segment $\gamma : [a, b] \to X$ with $\gamma(a) = x$ and $\gamma(b) = y$.*

(2) *For all points $x \in X$, $y \in \partial X$, there exists a k'-roughly geodesic ray $\gamma : [0, \infty) \to X$ with $\gamma(0) = x$ and $\lim_{t \to \infty} \gamma(t) = y$.*

(3) *For all points* $x, y \in \partial X$, $x \neq y$, *there exists a* k'-*rough geodesic* $\gamma : \mathbb{R} \to X$ *with* $\lim_{t \to -\infty} \gamma(t) = x$ *and* $\lim_{t \to \infty} \gamma(t) = y$.

Proof. All three statements have similar proofs, so we prove only (2). For each $t \geqslant 0$, let $z_t \in X$ be some point satisfying $(y|z_t) \geqslant t$ and let $y_t \in X$ satisfy $|y_t| \leqslant t + k$ and $|z_t - y_t| \leqslant |z_t| - t + k$. Similarly, let $x_t \in X$ satisfy $|x_t| \leqslant t + k$ and $|x - x_t| \leqslant |x| - t + k$. Let $a := (x|y)$, $b := |x|$, and set $\gamma(s) := x_{b-s}$ for $s \in [0, b - a]$ and $\gamma(s) := y_{2a-b+s}$ for $s > b - a$. It is left to the reader to verify that γ satisfies the requirements. $\square$

It follows from the definitions that each roughly geodesic metric space is also an almost geodesic metric space. Part (1) of the previous proposition implies the converse of this statement for Gromov hyperbolic spaces.

REMARK 5.3. The reader may wonder why we introduced the concepts of almost geodesic spaces, roughly geodesic rays, etc. It seems that less baggage of terminology would be needed if we had restricted ourselves to geodesic metric spaces. Even without being able to prove statements in the utmost generality, the essential features and ideas of the theory might be preserved.

Parts (2) and (3) of Proposition 5.2 can serve as a justification for introducing the concepts of roughly geodesic rays and rough geodesics. When we assume that X is geodesic, then the existence of geodesic rays and geodesics in parts (2) and (3) will not be true in general, unless we make additional assumptions, for example, that X is proper (closed balls are compact). In addition to that, the concepts of almost and roughly geodesic spaces, and the corresponding notions of roughly geodesic segments, etc., are better adapted to the philosophy of Gromov hyperbolic spaces: Bounded distortions do not matter, since they do not affect the structure of the space on large scales. While the property of being a geodesic metric space is very sensitive to these distortions, the property of being almost and roughly geodesic is more robust, and preserved under rough isometries, for example.

In spite of this, sometimes it does make sense to restrict oneself to geodesic metric spaces. A handy tool for generalizing results about Gromov hyperbolic geodesic metric spaces to the general case is Theorem 4.1. For example, geodesic stability generalizes in the following way.

PROPOSITION 5.4. *Let* X *be a* δ-*hyperbolic metric space, and* γ_1, γ_2 *be* (λ, k)-*roughly quasi-isometric paths in* X *with the same endpoints. Then* $d_{\text{Haus}}(\gamma_1, \gamma_2) = O_{\delta, \lambda, k}(1)$.

Proof. Embed X in a δ-hyperbolic geodesic metric space Y. If x and y are the common endpoints of γ_1, γ_2, and $[x, y]$ is any geodesic segment in Y joining x and y, then the geodesic stability of Y implies $d_{\text{Haus}}(\gamma_1, [x, y]) = O_{\delta,\lambda,k}(1)$ and $d_{\text{Haus}}(\gamma_2, [x, y]) = O_{\delta,\lambda,k}(1)$. The claim follows from the triangle inequality. □

PROPOSITION 5.5. *Suppose X and Y are δ-hyperbolic k-almost geodesic metric spaces, and $f : X \to Y$ is a (λ, c)-rough quasi-isometry of X into Y. Then for all $x, y, z, w \in X$*

(1) $\lambda^{-1}(x|y)_w - C \leqslant (x'|y')_{w'} \leqslant \lambda(x|y)_w + C,$

(2) *and if $(x|z)_w - (x|y)_w \geqslant 0$, then*

$$\lambda^{-1}\big((x|z)_w - (x|y)_w\big) - C \leqslant (x'|z')_{w'} - (x'|y')_{w'} \leqslant \lambda\big((x|z)_w - (x|y)_w\big) + C.$$

Here, image points under f are indicated by a prime, and $C = C(\delta, k, \lambda, c)$.

Statement (1) will be used to define the map $\partial f : \partial X \to \partial Y$ for a rough quasi-isometry $f : X \to Y$ (cf. Prop. 6.3). From Corollary 6.4, which is a generalization of (2), it will follow that ∂f is a power quasisymmetry (cf. Thm. 6.5).

Note that if $(x|z)_w - (x|y)_w \leqslant 0$ in (2), we get a similar inequality with the roles of λ and λ^{-1} reversed. To see this exchange y and z. This shows that (2) holds under the slightly weaker assumption $(x|z)_w - (x|y)_w \geqslant -C\delta$. We will need this improved version of (2) in the proof of Corollary 6.4.

Proof. The statement is essentially Prop. 15 on pp. 89–90 in [GH]. We will only outline the main idea of the proof. In order to prove (2) ((1) is just a special case putting $y = w$) we note that the expression $(x|z)_w - (x|y)_w$ is roughly the distance between the rough center of the triangle $[x, z, w]$ and the rough center of the triangle $[x, y, w]$. From Proposition 5.4 it follows that the image of the rough center of $[x, z, w]$ is within bounded distance to the rough center of $[x', z', w']$, and similarly for $[x, y, w]$. The proposition follows. □

A metric space X is called *k-visual with respect to $o \in X$*, if each point in X lies on a k-roughly geodesic ray emanating from o. We call X *k-visual*, if it is k-visual with respect to some point $o \in X$.

It is easy to see that if there exists a point $o \in X$ such that the union of the k_1-roughly geodesic rays from o is k_2-cobounded, then X is k-visual with $k = k(k_1, k_2)$. The following proposition gives an improved version of this statement for Gromov hyperbolic spaces.

PROPOSITION 5.6. *Let X be a δ-hyperbolic metric space with basepoint o. Suppose that the union of all (λ, c)-roughly quasi-isometric rays emanating*

from o is d-cobounded in X. Then X is k-visual with respect to o and k-roughly geodesic, where $k = k(\delta, \lambda, c, d)$.

The proposition implies in particular that every visual Gromov hyperbolic space is roughly geodesic.

Proof. Since the proof only employs standard techniques, we will only outline the main ideas and leave the details to the reader.

It is not hard to show that if $\gamma : [0, \infty) \to X$ is a (λ, c)-rough quasi-isometry into some δ-hyperbolic space, then there exists a map $s : [0, \infty) \to [0, \infty)$ with $s(0) = 0$ such that $\gamma' = \gamma \circ s$ is a $C(\delta, \lambda, c)$-rough isometry into X and $d_{\text{Haus}}(\gamma, \gamma') = O_{\delta, \lambda, c}(1)$. In other words, a roughly quasi-isometric ray admits a "rough reparameterization" to a roughly geodesic ray.

From this statement it follows that under the assumptions of the proposition the union of all k-roughly geodesic rays emanating from o is all of X for appropriate $k = k(\delta, \lambda, c, d)$.

In order to show the second statement let $x, y \in X$ be arbitrary. By the first part of the proof, there exist $C(\delta, \lambda, c, d)$-roughly geodesic rays γ_1, γ_2 emanating from o such that $x \in \gamma_1$ and $y \in \gamma_2$. Now a $C(\delta, \lambda, c, d)$-roughly geodesic segment joining x and y can be found by glueing together appropriate pieces of γ_1 and γ_2. □

6   The Boundary Functor ∂

The relation of a Gromov hyperbolic space to its (Gromov) boundary ∂X has been extensively studied (the definitions are given below). Suppose that $f : X \to Y$ is rough quasi-isometry of Gromov hyperbolic geodesic spaces. It is well known that there is an induced map $\partial f : \partial X \to \partial Y$, and ∂f is Hölder and quasiconformal. Indeed ∂f satisfies a stronger condition, namely, it is a power quasisymmetry. This observation is used to show that under certain mild conditions X is determined by the quasisymmetry class of ∂X up to rough quasi-isometries. F. Paulin [P] obtained a different characterization of X from ∂X, but it seems that for his characterization one must assume that there is a group acting co-compactly on X.

The next goal is to identify the precise relevant structure on ∂X that is induced by the metric on X. The correct structure would depend on how seriously the metric on X is taken. More precisely, if you are interested only in properties of X that are invariant under rough quasi-isometries, the corresponding structure on ∂X would be coarser than if you were interested in X up to rough isometries.

DEFINITION.   Let X and Y be two metric spaces, $f : X \to Y$ be a bijection and $\alpha > 0$, $\lambda \geqslant 1$.

(1) The map f is an (α, λ)-*snowflake map* if for all $x, y \in X$
$$\lambda^{-1} |x - y|^\alpha \leqslant \big| f(x) - f(y) \big| \leqslant \lambda |x - y|^\alpha \,.$$

(2) The map f is an (α, λ)-*quasisymmetry* if for all distinct points $x, y, z \in X$
$$\frac{|f(x) - f(z)|}{|f(x) - f(y)|} \leqslant \eta_{\alpha,\lambda} \left(\frac{|x - z|}{|x - y|} \right) .$$

Here
$$\eta_{\alpha,\lambda}(t) = \begin{cases} \lambda t^{1/\alpha} & \text{for} \quad 0 < t < 1 \,, \\ \lambda t^\alpha & \text{for} \quad 1 \leqslant t \,. \end{cases}$$

It seems that the class of maps in (1) has not been given a name in the literature. We call them snowflake maps, because these maps behave similar as the map giving the parameterization of the well-known von Koch snowflake curve.

In contrast to this, quasisymmetries have appeared in the literature before (cf. [TuV1], [V]). Our usage of this term slightly differs from the common one. In general, a map is called a quasisymmetry if the above condition holds for some increasing homeomorphism $\eta : (0, \infty) \to (0, \infty)$. So our notion of quasisymmetry is stronger. To emphasize this distinction we call a map a *power quasisymmetry*, if it is an (α, λ)-quasisymmetries for some $\alpha > 0$ and $\lambda \geqslant 1$. For some spaces, e.g., connected spaces, quasisymmetries are always power quasisymmetries (cf. [TuV1, Cor. 3.12]).

Every snowflake map is a power quasisymmetry. It is straightforward to check that inverse maps and compositions of snowflake maps are again snowflake maps. The same statement is true for power quasisymmetries. In fact, the inverse of an (α, λ)-quasisymmetry is an $(\alpha, C(\lambda, \alpha))$-quasisymmetry. Quasisymmetries are homeomorphisms.

Two metrics d_1 and d_2 on a space X are called *bilipschitz equivalent*, *snowflake equivalent* or *quasisymmetrically equivalent* if the identity map $\mathrm{id}_X : (X, d_1) \to (X, d_2)$ is bilipschitz, a snowflake map or a quasisymmetry, respectively. This defines equivalence relations for the metrics on a space X.

Following is the standard construction of the metrics on ∂X, where X is a Gromov hyperbolic space. If $x, y \in \partial X, w \in X, \epsilon > 0$, let

$$d_{\partial X, w, \epsilon}(x, y) = d_{w, \epsilon}(x, y) = \inf \left\{ \sum_{i=1}^{n} e^{-\epsilon (x_{i-1} | x_i)_w} \right\}, \qquad (6.1)$$

where the infimum extends over all finite sequences $x = x_0, x_1, x_2, \ldots, x_n = y$ in ∂X. Here the convention $e^{-\infty} = 0$ is understood.

LEMMA 6.1. *There is some constant $\epsilon_0 > 0$ with the following property. If X is δ-hyperbolic and $\epsilon\delta \leqslant \epsilon_0$, then*

$$\tfrac{1}{2}e^{-\epsilon(x|y)_w} \leqslant d_{w,\epsilon}(x,y) \leqslant e^{-\epsilon(x|y)_w}, \quad \forall\, x,y \in \partial X. \tag{6.2}$$

For the proof, see, for example, [GH, Ch. 7].

This lemma leads to the following definitions.

DEFINITIONS. The *canonical gauge* $\mathcal{G}(X)$ on ∂X is the set of all metrics of the form $d = d_{w,\epsilon}$. We say that (ϵ_1, d_1) is *B-equivalent* to (ϵ_2, d_2) if there is a constant $c > 0$ such that

$$c^{-1}d_2^{\epsilon_1} \leqslant d_1^{\epsilon_2} \leqslant c\, d_2^{\epsilon_1}.$$

A *B-structure* on Z is an equivalence class of this equivalence relation.

The lemma tells us that $(\epsilon, d_{w,\epsilon})$ and $(\epsilon', d_{w',\epsilon'})$ are B-equivalent when $w, w' \in X$ and $\epsilon, \epsilon' > 0$ satisfy $\epsilon\delta \leqslant \epsilon_0$ and $\epsilon'\delta \leqslant \epsilon_0$. Therefore, when X is Gromov hyperbolic, there is an associated B-structure on ∂X. It is called the *canonical B-structure* on ∂X.

Mostly, it is convenient to work with a fixed metric in the canonical gauge and to suppress the ambiguity in choosing this metric. We often abbreviate such a metric by $d_{\partial X}$.

PROPOSITION 6.2. *The boundary ∂X of any Gromov hyperbolic space X is bounded and complete.*

Note that boundedness and completeness are properties that do not depend on the particular choice of the metric $d_{\partial X}$ in $\mathcal{G}(X)$. There are situations where ∂X is not compact. For example, when X is infinite dimensional hyperbolic space, ∂X is homeomorphic to the unit sphere in Hilbert space.

Proof. The boundedness of ∂X follows from the fact that the Gromov product is nonnegative.

Fix a basepoint $o \in X$. Let $\{y_i\} \subset \partial X$ be a $d_{\partial X}$-Cauchy sequence for some metric in the canonical gauge. Then $\lim_{i,j\to\infty}(y_i|y_j) = \infty$. There exists a sequence $\{x_i\} \subset X$ with $\lim_{i\to\infty}(x_i|y_i) = \infty$. Then $\{x_i\}$ converges at infinity. If $y \in \partial X$ is the equivalence class of $\{x_i\}$, then $\lim_{i\to\infty} y_i = y$. $\square$

We will now discuss the functorial properties of $X \mapsto \partial X$.

PROPOSITION 6.3. *Suppose X, Y, Z are Gromov hyperbolic almost geodesic metric spaces, and $f : X \to Y$, $g : Y \to Z$ are roughly quasi-isometric embeddings.*

(1) *If $\{x_i\} \subset X$ converges at infinity, then $\{f(x_i)\} \subset Y$ converges at infinity. If $\{x_i\}$ and $\{y_i\}$ are equivalent sequences in X converging at infinity, then $\{f(x_i)\}$ and $\{f(y_i)\}$ are also equivalent.*

(2) If $a \in \partial X$ and $\{x_i\} \in a$, let $b \in \partial Y$ be the equivalence class of $\{f(x_i)\}$ and define $\partial f(a) = b$. Then $\partial f : \partial X \to \partial Y$ is well-defined. Moreover, $\partial(g \circ f) = \partial g \circ \partial f$.

(3) If $f_1, f_2 : X \to Y$ are rough quasi-isometric embeddings and $f_1 \approx f_2$, then $\partial f_1 = \partial f_2$.

(4) The map ∂f is injective. If the image of f is cobounded in Y, i.e., if f is a rough quasi-isometry, then ∂f is a bijection.

Proof. The statements in (1)–(3) immediately follow from the definitions and (1) of Proposition 5.5. For the first part of (4) note that if $\{x_i\}$ and $\{y_i\}$ are sequences in X converging at infinity such that $\{f(x_i)\}$ and $\{f(y_i)\}$ are equivalent, then $\lim_{i\to\infty}(f(x_i)|f(y_i)) = \infty$. This implies $\lim_{i\to\infty}(x_i|y_i) = \infty$ by 5.5.(2). Thus $\{x_i\}$ and $\{y_i\}$ are equivalent.

If f is a rough quasi-isometry, then it has a rough inverse $h : Y \to X$ which is a rough quasi-isometry. Since $\partial(\mathrm{id}_X) = \mathrm{id}_{\partial X}$, the statements (2), (3) and the relations $h \circ f \approx \mathrm{id}_X$, $f \circ h \approx \mathrm{id}_Y$ imply $\partial h \circ \partial f = \mathrm{id}_{\partial X}$ and $\partial f \circ \partial h = \mathrm{id}_{\partial Y}$. From this it follows that ∂f is bijective. $\square$

COROLLARY 6.4.   *The inequalities in Proposition 5.5 are also valid if $x, y, z \in \partial X$, if x', y', z' are the images of these points under ∂f, and if we assume in part (2) that x, y, z are distinct points.*

Proof. This follows from the definition of ∂f, Proposition 5.5, and (3.4). Note that for the proof of the second part we need inequality 5.5.(2) under the slightly stronger assumption $(x|z)_w - (x|y)_w \geqslant -C\delta$, $x, y, z, w \in X$ (cf. the remark following Prop. 5.5). $\square$

Theorem 6.5. *Suppose X and Y are Gromov hyperbolic almost geodesic metric spaces, and let $f : X \to Y$ be a map.*

(1) *If f is a rough similarity, then ∂f is a snowflake map.*

(2) *If f is a rough quasi-isometry, then ∂f is a power quasisymmetry.*

Proof. Let x' denote $\partial f(x)$ for $x \in \partial X$. Note that ∂f is bijective by Proposition 6.3 (4).

The properties of ∂f in question do not depend on which metrics $d_{\partial X} \in \mathcal{G}(X)$ and $d_{\partial Y} \in \mathcal{G}(Y)$ in the canonical gauges on ∂X and ∂Y we choose. We may assume that the image w' of the basepoint w of X is the basepoint in Y. Then there exists constant $\epsilon, \epsilon' > 0$ such that for all $x, y \in \partial X$

$$d_{\partial X}(x, y) \asymp e^{-\epsilon(x|y)} \quad \text{and} \quad d_{\partial Y}(x', y') \asymp e^{-\epsilon'(x'|y')}. \qquad (6.3)$$

Here $\asymp$ means equality up to a multiplicative constant independent of x and y.

From this we see that statement (1) is equivalent to the existence of constants $\lambda > 0$ and $K \geqslant 0$ such that
$$\lambda(x|y) - K \leqslant (x'|y') \leqslant \lambda(x|y) + K \tag{6.4}$$
for all $x, y \in \partial X$. If f is a rough similarity, an inequality like this holds for $x, y \in X$ (the prime indicating the image under f in this case). The definition of ∂f, and (3.4) imply that (6.4) also holds for $x, y \in \partial X$.

Statement (2) is a straightforward consequence of (6.3), 5.5.(2), and Corollary 6.4. $\square$

We can summarize Proposition 6.2, Proposition 6.3, and Theorem 6.5 as follows. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be categories, where the objects are Gromov hyperbolic almost geodesic metric spaces. The morphisms of $\mathcal{C}_1$ are rough similarities, while the morphisms of $\mathcal{C}_2$ are rough quasi-isometries. Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be categories, where the objects are bounded and complete B-structures. The morphisms of $\mathcal{D}_1$ are snowflake maps, and the morphisms of $\mathcal{D}_2$ are power quasisymmetries. Then $X \mapsto \partial X$ is a functor from $\mathcal{C}_i$ to $\mathcal{D}_i$ for $i \in \{1, 2\}$. Actually, as the results in section 7 will indicate, it is more appropriate to consider rough mapping classes of rough similarities and rough quasi-isometries as the morphisms of $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively.

7   The Metric Space Con(Z)

Given a bounded metric space (Z, d), we now construct a Gromov hyperbolic space $\mathrm{Con}(Z)$. The space $\mathrm{Con}(Z)$ has properties analogous to the hyperbolic convex hull of a set in the boundary of a real hyperbolic space. We refer to section 10 for some further discussion. Our construction is similar to one given by Gromov [Gr, 1.8.A.(b)] and to a construction of Trotsenko and Väisälä [TV].

Set
$$\mathrm{Con}(Z) = Z \times \big(0, D(Z)\big].$$
Here and in the following we let $D(Z) = \mathrm{diam}(Z)$ if $\mathrm{diam}(Z) > 0$, and $D(Z) = 1$ if $\mathrm{diam}(Z) = 0$, i.e., if Z consists of a single point. It is convenient to include this trivial case in the definition, for otherwise we would have to exclude it explicitly in some of the following statements.

Define $\rho : \mathrm{Con}(Z) \times \mathrm{Con}(Z) \to [0, \infty)$ by
$$\rho\big((z, h), (z', h')\big) = 2 \log \left(\frac{d(z, z') + h \vee h'}{\sqrt{hh'}} \right). \tag{7.1}$$

The motivation of the factor 2 comes from the fact that it corresponds to curvature -1 for real hyperbolic spaces.

LEMMA 7.1. ρ is a metric on $\mathrm{Con}(Z)$.

Proof. The triangle inequality,

$$\rho\big((z,h),(z'',h'')\big) \leqslant \rho\big((z,h),(z',h')\big) + \rho\big((z',h'),(z'',h'')\big), \quad (7.2)$$

reduces to

$$\frac{d(z,z'') + h \vee h''}{\sqrt{hh''}} \leqslant \frac{d(z,z') + h \vee h'}{\sqrt{hh'}} \cdot \frac{d(z',z'') + h' \vee h''}{\sqrt{h'h''}},$$

which is equivalent to

$$h'\big(d(z,z'') + h \vee h''\big) \leqslant \big(d(z,z') + h \vee h'\big)\big(d(z',z'') + h' \vee h''\big). \quad (7.3)$$

Since $d(z,z'') \leqslant d(z,z') + d(z',z'')$, inequality (7.3) holds, which verifies the triangle inequality. The other properties of a metric are immediate. $\square$

Henceforth, $\mathrm{Con}(Z)$ will always be equipped with this metric ρ. We will use the same letter ρ for metrics on different spaces $\mathrm{Con}(Z)$.

Theorem 7.2.   *There are constants $\delta \geqslant 0$, $k \geqslant 0$ with the following property. If (Z,d) is a bounded metric space, then $\mathrm{Con}(Z)$ is δ-hyperbolic, k-visual, and k-roughly geodesic.*

Proof. First we prove the Gromov hyperbolicity of $\mathrm{Con}(Z)$.

Suppose we are given numbers $r_{ij} \geqslant 0$ such that $r_{ij} = r_{ji}$ and $r_{ij} \leqslant r_{ik} + r_{kj}$ for $i,j,k \in \{1,2,3,4\}$. Then $r_{12}r_{34} \leqslant 4((r_{13}r_{24}) \vee (r_{14}r_{23}))$. To see this, we may assume that r_{13} is the smallest of the quantities r_{ij} appearing on the right hand side of this inequality. Then $r_{12} \leqslant r_{13} + r_{32} \leqslant 2r_{23}$ and $r_{34} \leqslant r_{31} + r_{14} \leqslant 2r_{14}$. The inequality follows.

Now let $x_i = (z_i, h_i)$, $i \in \{1,2,3,4\}$, be four arbitrary points in $\mathrm{Con}(Z)$. Set $d_{ij} = d(z_i, z_j)$ and $r_{ij} = d_{ij} + h_i \vee h_j$. The numbers r_{ij} satisfy the above requirements. Hence

$$(d_{1,2} + h_1 \vee h_2)(d_{3,4} + h_3 \vee h_4)$$
$$\leqslant 4\big((d_{1,3} + h_1 \vee h_3)(d_{2,4} + h_2 \vee h_4)\big) \vee \big((d_{1,4} + h_1 \vee h_4)(d_{2,3} + h_2 \vee h_3)\big).$$

This translates to

$$\rho(x_1, x_2) + \rho(x_3, x_4) \leqslant \big(\rho(x_1, x_3) + \rho(x_2, x_4)\big) \vee \big(\rho(x_1, x_4) + \rho(x_2, x_3)\big) + C,$$

which shows that $\mathrm{Con}(Z)$ is C-hyperbolic (cf. (3.3)).

Choose a point $z_0 \in Z$, and let $o = (z_0, D(Z)) \in \mathrm{Con}(Z)$. If $p = (z,h) \in \mathrm{Con}(Z)$ is an arbitrary point, a C-roughly geodesic ray emanating from o and passing through p can be obtained from an appropriate parameterization of the set $\{o\} \cup \{z\} \times (0, D(Z)]$. More precisely, define the ray $\gamma : [0, \infty) \to \mathrm{Con}(Z)$ by $\gamma(0) = o$ and $\gamma(t) = (z, D(Z)e^{-t})$ for $t > 0$. It is easy to check that γ has the required properties.

Finally, $\mathrm{Con}(Z)$ is C-roughly geodesic as follows from Proposition 5.6 and the first two parts of the proof.   □

We want to show that the assignment $Z \mapsto \mathrm{Con}(Z)$ is a functor of appropriately defined categories. The basic problem is to define an induced map $\widehat{f} : \mathrm{Con}(X) \to \mathrm{Con}(Y)$ for a map $f : X \to Y$, $x \mapsto x'$. This is possible if f is a power quasisymmetry. As we mentioned in the introduction, the idea is similar to the fuzzy extension of a quasiconformal map on $\mathbb{R}^n$ to upper half-space in $\mathbb{R}^{n+1}$ by Tukia and Väisälä [TuV2]. One would like to define $\widehat{f}(x, h) = (x', |x' - y'|)$, where y is a point in X with $h = |x - y|$. The fact that $\widehat{f}$ may not be well-defined is not serious problem. For if z is any other point with $h = |x - z|$, then $|x' - y'|$ and $|x' - z'|$ are comparable if f is a power quasisymmetry. In this case, the ρ-distance of $(x', |x' - y'|)$ and $(x', |x' - z'|)$ in $\mathrm{Con}(Y)$ is bounded by a constant only depending on the parameters of f. So $\widehat{f}$ is "roughly" well-defined.

More serious is the problem that there might exist no y with $h = |x - y|$, or even worse, no y such that $|x - y|$ is comparable to h. In this case, one has to "interpolate" between those heights h which are attained as distances of points.

The following lemma gives the basic step in the construction of $\widehat{f}$. For $z \in Z$, we denote by R_z the ray in $\mathrm{Con}(Z)$ that ends at $z \in Z$: $R_z = \{z\} \times (0, D(Z)] \subset \mathrm{Con}(Z)$. The maps f_x of the lemma will be used in the definition of $\widehat{f}$, and will be the restrictions of $\widehat{f}$ to the rays R_x (cf. the definition following Lemma 7.3).

LEMMA 7.3.   Suppose that $f : X \to Y$, $x \mapsto x' = f(x)$, is an (α, c)-quasisymmetry of bounded metric spaces X, Y. For all $x \in X$ we can define a map $f_x : R_x \to R_{x'}$ with the following properties:

(1) There exist $\lambda = \lambda(\alpha, c) \geq 1$, $k = k(\alpha, c) \geq 0$ such that all maps f_x, $x \in X$, are (λ, k)-rough quasi-isometries.

(2) If $x, y \in X$, $x \neq y$, $h = |x - y|$, and $f_x(x, h) = (x', h')$, then

$$C(\alpha, c)^{-1} |x' - y'| \leq h' \leq |x' - y'| C(\alpha, c).$$

(3) Let $x \in X$, $t_1, t_2 \in (0, D(X)]$, $f_x(x, t_1) = (x', t_1')$, $f_x(x, t_2) = (x', t_2')$. Then

$$t_1 \leq t_2 \Rightarrow t_1' \leq t_2',$$

     i.e., f_x is non-decreasing on R_x.

(4) For all $x, y \in X$, $x \neq y$, and $h \in [|x - y|, D(X)]$

$$\rho\big(f_x(x, h), f_y(y, h)\big) = O_{\alpha, c}(1).$$

(5) If Z is a metric space with finite diameter and $g : Y \to Z$ is a (β, c')-quasisymmetry, then for all $x \in X$ and $p \in R_x$

$$\rho\big((g \circ f)_x(p), g_{x'}(f_x(p))\big) = O_{\alpha,\beta,c,c'}(1) \, .$$

In particular, $(g \circ f)_x \approx g_{f(x)} \circ f_x$.

(6) If f is a bilipschitz or snowflake map, then the maps f_x, $x \in X$, can be defined such that they are rough isometries or rough similarities with the same parameters, respectively. The parameters only depend on the parameters of f. These maps satisfy statements similar to (2)–(5) (cf. proof).

Proof. We may assume $\mathrm{diam}(X) > 0$. For $x \in X$ let $S_x \subset \mathbb{N}$ be the set of all $l \in \mathbb{N}$ for which the annulus $A_X(x, l) = \{z \in X : 2^{-l-1}\,\mathrm{diam}(X) < |z - x| \leqslant 2^{-l}\,\mathrm{diam}(X)\}$ is nonempty. The set S_x can be regarded as the "scale spectrum" of X at x. Note that $m = \min S_x \in \{0, 1\}$. Similarly, let $S_{x'}$ be the scale spectrum of Y at the point x'.

We define a map $\phi_x : S_x \to S_{x'}$ by

$$\phi_x(l) = \sup\left\{l' \mid \exists y \in X : 2^{-l-1}\,\mathrm{diam}(X) < |y - x|\,, \ y' \in A_Y(x', l')\right\}, \quad \forall l \in S_x \, .$$

Obviously, ϕ_x is non-decreasing on S_x. If $l \in S_x$ there exists a point $y \in A_X(x, l)$. If we take any such point y, then for a unique $l' \in S_{x'}$ we have $y' \in A_Y(x', l')$. From the fact that f is an (α, c)-quasisymmetry it follows that $|l' - \phi_x(l)| = O_{\alpha,c}(1)$.

For $l_1, l_2 \in S_x$ choose points $y_1 \in A_X(x, l_1)$, and $y_2 \in A_X(x, l_2)$. Since f is an (α, c)-quasisymmetry,

$$C(\alpha, c)^{-1} \eta_{\alpha,c}(2^{l_1-l_2})^{-1} \leqslant \frac{|y_1' - x'|}{|y_2' - x'|} \leqslant C(\alpha, c)\eta_{\alpha,c}(2^{l_2-l_1}) \, . \quad (7.4)$$

This and the remark following the definition of ϕ_x show

$$\lambda^{-1}|l_2 - l_1| - C(\alpha, c) \leqslant |\phi_x(l_2) - \phi_x(l_1)| \leqslant \lambda|l_2 - l_1| + C(\alpha, c) \, , \quad (7.5)$$

where $\lambda = \lambda(\alpha) \geqslant 1$.

The set $\phi_x(S_x)$ is $C(\alpha, c)$-cobounded in $S_{x'}$. For if $l' \in S_{x'}$ is arbitrary, then, since f is bijective, there exists a point $y \in X$, $y \neq x$, such that $y' \in A_Y(x', l')$. For some $l \in S_x$ we have $y \in A_X(x, l)$. By the above remark $|\phi_x(l) - l'| = O_{\alpha,c}(1)$.

Since $\{0, 1\} \cap S_{x'} \neq 0$ and ϕ_x is non-decreasing, this coboundedness implies

$$\phi_x(m) \leqslant C(\alpha, c) \, . \quad (7.6)$$

We now extend the map $\phi_x : S_x \to S_{x'}$ to a map $\phi_x : [0, \infty) \to [0, \infty)$. We use the same notation for the new and the old map. The extension is essentially obtained by linear interpolation.

More precisely, let $M = \sup S_x \in \mathbb{N} \cup \{\infty\}$. Define $\phi_x(t) = \phi_x(m)$ for $t \in [0, m)$. If $t \in (m, M) \backslash S_x$, there a smallest interval (l_1, l_2) with endpoints $l_1, l_2 \in S_x$ containing t. Then $t = \mu l_1 + (1 - \mu) l_2$ for a unique $\mu \in (0, 1)$. Define $\phi_x(t) = \mu \phi_x(l_1) + (1 - \mu) \phi_x(l_2)$. Finally, if $M < \infty$ and $t > M$ put $\phi_x(t) = \phi_x(M) + (t - M)$.

Obviously, ϕ_x is continuous and non-decreasing. Furthermore, $\lim_{t \to \infty} \phi_x(t) = \infty$. This follows from the definition of ϕ_x if $M < \infty$ and from (7.5) if $M = \infty$. Now (7.6) shows that $\phi_x([0, \infty))$ is $C(\alpha, c)$-cobounded in $[0, \infty)$.

Since $\phi_x : S_x \to S_{x'}$ is non-decreasing and satisfies (7.5), it follows that $\phi_x : [0, \infty) \to [0, \infty)$ satisfies (7.5) for all $l_1, l_2 \in [0, \infty)$. This statement is straightforward to prove. We leave the details to the reader. (Note that (7.5) is true, when l_1, l_2 are in some minimal non-degenerate interval with endpoints in $S_x \cup \{0, \infty\}$. Introducing appropriately chosen points in $S_x \cap [l_1, l_2]$, the general case of (7.5) can be reduced to this special case.)

These considerations show that $\phi_x : [0, \infty) \to [0, \infty)$ is a non-decreasing rough quasi-isometry with parameters only depending on α and c.

Now let $D = \operatorname{diam}(X)$, $D' = \operatorname{diam}(Y)$, and define for $(x, h) \in \operatorname{Con}(X)$

$$f_x(x, h) = \left(x', D' 2^{-\phi_x(\log_2(D/h))}\right). \tag{7.7}$$

The statements about the maps f_x follow directly from the properties of the maps ϕ_x. This is clear for (1) and (3). For (2) note that $\phi_x(\log_2(D/|x - y|)) = \log_2(D'/|x' - y'|) + O_{\alpha,c}(1)$ for $x, y \in X$, $x \neq y$. Statement (4) follows from (2) and (3).

To prove (5), define $\psi_y : S_y \to S_{g(y)}$, $y \in Y$, and $\theta_x : S_x \to S_{g(f(x))}$, $x \in X$, in the same way for g and $g \circ f$, respectively, as the maps ϕ_x, $x \in X$, were defined for f. Let $x \in X$. Making the parameters of the maps f and g larger if necessary, we may assume that both are power quasisymmetries with the same parameters, α and c, say. This implies that for $t \in S_x$

$$\theta_x(t) = \psi_{x'}(\phi_x(t)) + O_{\alpha,c}(1). \tag{7.8}$$

To establish this equation for general $t \in [0, \infty)$ assume $t \in (m, M) \backslash S_x$. The cases $t \in [0, m]$ or $t \in [M, \infty)$ are easier and can be treated in a similar way as below.

Let (l_1, l_2) be the smallest interval with endpoints in S_x containing x and write $t = \mu l_1 + (1 - \mu) l_2$ with $\mu \in (0, 1)$. Then $\phi_x(t) = \mu l'_1 + (1 - \mu) l'_2$, where $l'_1 = \phi_x(l_1)$ and $l'_2 = \phi_x(l_2)$. In particular, $\phi_x(t) \in [\min S_{x'}, \sup S_{x'}]$, and so there exists some minimal interval $[\tilde{l}_1, \tilde{l}_2]$ with endpoints in $S_{x'}$ containing $\phi_x(t)$. The interval $[\tilde{l}_1, \tilde{l}_2]$ is degenerate if $\phi_x(t) \in S_{x'}$.

Since ϕ_x is non-decreasing and $\phi(S_x)$ is $C(\alpha, c)$-cobounded in $S_{x'}$, we have $\tilde{l}_1 = l'_1 + O_{\alpha,c}(1)$ and $\tilde{l}_2 = l'_2 + O_{\alpha,c}(1)$. From the definition of θ_x and ϕ_x, equation (7.8) for $t \in S_x$, and the fact that ψ_x is a rough quasi-isometry with parameters only depending on α and c, we obtain

$$\theta_x(t) = \mu\psi_{x'}(l'_1) + (1 - \mu)\psi_{x'}(l'_2) + O_{\alpha,c}(1)$$
$$= \mu\psi_{x'}(\tilde{l}_1) + (1 - \mu)\psi_{x'}(\tilde{l}_2) + O_{\alpha,c}(1)$$
$$= \psi_{x'}(\phi_x(t)) + O_{\alpha,c}(1) \,.$$

This shows that (7.8) holds for all $t \in [0, \infty)$. The maps $(g \circ f)_x$, $x \in X$, and g_y, $y \in Y$, are defined similarly as in (7.7) using the maps θ_x and ψ_y, respectively. Statement (5) then follows from (7.8).

To see that (6) is true, define $f_x(x, h) = (x', (D'/D)h)$ for $(x, h) \in \mathrm{Con}(X)$ if f is a bilipschitz map. Statements (1)–(5) are immediate in this case with all constants only depending on the bilipschitz constant of f (and the bilipschitz constant of g in (5)).

Suppose that f is an (α, c)-snowflake map. If $S(X) := \bigcup_{x \in X} S_x$ is an infinite set, then α is uniquely determined. In this case let $\alpha(f) = \alpha$. If $S(X)$ is a finite set, let $\alpha(f) = 1$.

Note that the property of a space X that $S(X)$ is finite or infinite is invariant under snowflake maps. This implies that if $g : Y \to Z$ is a snowflake map, then $\alpha(g \circ f) = \alpha(g)\alpha(f)$.

The set $S(X)$ is finite, if and only if $d(X) := \inf_{x,y \in X, x \neq y} |x - y| > 0$. In this case, f is a λ-bilipschitz map with $\lambda = \lambda(\alpha, c, \mathrm{diam}(X), d(X))$.

Now define $f_x(x, h) = (x', D'(h/D)^{\alpha(f)})$.

The stated properties of the maps f_x are immediate to check. The constants in (1), (2), (4) will only depend on α and c. If $S(X)$ is finite, the constant in (2) will also depend on $d(X)$ and $\mathrm{diam}(X)$. The snowflake version of statement (5) is true with constant 0.   □

If X and Y are bounded metric spaces, and $f : X \to Y$ is a power quasisymmetry (which includes snowflake or bilipschitz maps), we define a map $\widehat{f} : \mathrm{Con}(X) \to \mathrm{Con}(Y)$ by

$$\widehat{f}(x, h) = f_x(x, h) \,, \quad \forall (x, h) \in \mathrm{Con}(X) \,.$$

Here f_x, $x \in X$, are the maps defined in the last lemma.

Theorem 7.4.   *Suppose that $f : X \to Y$ is a power quasisymmetry of bounded metric spaces X, Y. Then $\widehat{f} : \mathrm{Con}(X) \to \mathrm{Con}(Y)$ is a rough quasi-isometry. If f is a snowflake or bilipschitz map, then $\widehat{f}$ is a rough similarity or a rough isometry, respectively.*

If Z is a bounded metric space, and $g : Y \to Z$ is a map of the same type as f, i.e., a power quasisymmetry, a snowflake map, or a bilipschitz map, then $\widehat{g \circ f} \approx \widehat{g} \circ \widehat{f}$.

Proof. Assume $f : X \to Y$, $x \mapsto x' = f(x)$ is an (α, c)-quasisymmetry.

We have $\widehat{f}(\mathrm{Con}(X)) = \bigcup_{x \in X} f_x(R_x)$, and $\bigcup_{x \in X} R_{x'} = \mathrm{Con}(Y)$. Since $f_x(R_x)$ is $C(\alpha, c)$-cobounded in $R_{x'}$ for all $x \in X$ by Lemma 7.3.(1), the set $\widehat{f}(\mathrm{Con}(X))$ is $C(\alpha, c)$-cobounded in $\mathrm{Con}(Y)$.

To show that $\widehat{f} : \mathrm{Con}(X) \to \mathrm{Con}(Y)$ is a rough quasi-isometry, we have to establish inequality (2.1) for $\widehat{f}$. For two arbitrary points $q_1 = (x_1, h_1), q_2 = (x_2, h_2) \in \mathrm{Con}(X)$ we express the distance of q_1 and q_2 and the distance of their image points $q_1' = (x_1', h_1')$ and $q_2' = (x_2', h_2')$ under $\widehat{f}$ by distances only involving pairs of points lying on the same ray R_z. Then (2.1) will follow from the definition of $\widehat{f}$ and the fact that the maps f_x, $x \in X$, are rough quasi-isometries with the same parameters.

To that purpose define $h = |x_1 - x_2| \vee h_1 \vee h_2$, $p_1 = (x_1, h)$, $p_2 = (x_2, h)$, $p_1' = \widehat{f}(p_1) = (x_1', h')$, $p_2' = \widehat{f}(p_2) = (x_2', \tilde{h}')$. The definition of ρ shows

$$\rho(q_1, q_2) = \rho(q_1, p_1) + \rho(p_2, q_2) + O(1). \qquad (7.9)$$

If $h = |x_1 - x_2|$, then from Lemma 7.3.(2) and (3) it follows that

$$C(\alpha, c)^{-1}|x_1' - x_2'| \leqslant h', \tilde{h}' \leqslant C(\alpha, c)|x_1' - x_2'|, \quad \text{and} \quad h_1' \vee h_2' \leqslant h', \tilde{h}'.$$

This implies

$$\rho(q_1', q_2') = \rho(q_1', p_1') + \rho(p_2', q_2') + O_{\alpha, c}(1). \qquad (7.10)$$

If $h \neq |x_1 - x_2|$, i.e., if $h = h_1 \vee h_2$, we may without loss of generality assume that $h = h_1$. Then $q_1' = p_1'$ and Lemma 7.3.(4) shows that $\rho(p_1', p_2') = O_{\alpha, c}(1)$. Thus, (7.10) also holds in this case.

Consequently, Lemma 7.3.(1) now gives

$$\begin{aligned}
\rho(q_1', q_2') &= \rho(q_1', p_1') + \rho(p_2', q_2') + O_{\alpha, c}(1) \\
&\leqslant C(\alpha, c)\big(\rho(q_1, p_1) + \rho(p_2, q_2)\big) + O_{\alpha, c}(1) \\
&= C(\alpha, c)\rho(q_1, q_2) + O_{\alpha, c}(1).
\end{aligned}$$

An inequality in the opposite direction can be obtained in the same way. This shows that $\widehat{f}$ is a rough quasi-isometry.

If f is a snowflake or a bilipschitz map, the same argument based on (7.9), (7.10) and Lemma 7.3.(6) shows that $\widehat{f}$ is a rough similarity or a rough isometry, respectively. Finally, the last statement follows from Lemma 7.3.(5). $\qquad \square$

Again we may summarize the results of this section in the language of categories. Let $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ be categories, where the objects are bounded

metric spaces, and the morphisms are power quasisymmetries, snowflake maps, and bilipschitz maps, respectively.

Let $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$ be categories, where the objects are visual Gromov hyperbolic spaces, and the morphisms are rough mapping classes of rough quasi-isometries, rough similarities, and rough isometries, respectively. Then $Z \mapsto \mathrm{Con}(Z)$ is a functor from $\mathcal{C}_i$ to $\mathcal{D}_i$ for $i \in \{1, 2, 3\}$.

In particular, we get the following correspondence for the types of maps

$$
\begin{aligned}
\text{power quasisymmetry} \quad &\widehat{=} \quad \text{rough quasi-isometry},\\
\text{snowflake map} \quad &\widehat{=} \quad \text{rough similarity},\\
\text{bilipschitz map} \quad &\widehat{=} \quad \text{rough isometry}.
\end{aligned}
$$

Note that the first two correspondences (going from right to left) are exactly what we found in the last section for the functor $X \mapsto \partial X$.

8   The Relation of ∂X and $\mathrm{Con}(Z)$

In this section we will investigate the relation of the functors $X \mapsto \partial X$ and $Z \mapsto \mathrm{Con}(Z)$. We content ourselves with proving two statements about the objects of the categories involved. Similar questions can also be studied for the morphisms.

Theorem 8.1. *Suppose (Z, d) is a complete bounded metric space, and let $X = \mathrm{Con}(Z)$. Then the spaces ∂X and Z can be identified as sets, and d is bilipschitz to a metric in $\mathcal{G}(X)$, the canonical gauge on ∂X.*

Proof. Take as the basepoint $o \in X$ some point of the form (z_0, D), where $D = D(Z)$. Let $x = (z, h)$, $x' = (z', h') \in X$. Then

$$
\begin{aligned}
(x|x') &= \log\left(\frac{d(z_0, z) + D}{\sqrt{Dh}}\right) + \log\left(\frac{d(z_0, z') + D}{\sqrt{Dh'}}\right) - \log\left(\frac{d(z, z') + h \vee h'}{\sqrt{hh'}}\right)\\
&= -\log\left(d(z, z') + h \vee h'\right) + O_D(1).
\end{aligned}
$$

From this follows that a sequence $\{(z_i, h_i)\}$ in $\mathrm{Con}(Z)$ converges at infinity if and only if $\{z_i\}$ is a Cauchy sequence in Z and $\lim_{i \to \infty} h_i = 0$. Since Z is complete, $\{z_i\}$ has a limit y in Z. It is immediate to check that if a sequence $\{(z_i', h_i')\}$ is equivalent to $\{(z_i, h_i)\}$, then $\lim_{i \to \infty} z_i = \lim_{i \to \infty} z_i' = y$. So by assigning y to the equivalence class of $\{(z_i, h_i)\}$ we get a well-defined map from ∂X to Z. It is straightforward to verify that this map is bijective. Thus, ∂X and Z can be identified as sets.

Moreover, from the above expression and (3.4) we conclude

$$
(z|z') = -\log(d(z, z')) + O_D(1), \qquad \forall z, z' \in \partial X = Z.
$$

This shows that d is bilipschitz to a metric in the canonical gauge on ∂X. $\square$

Theorem 8.2. *Suppose X is a visual Gromov hyperbolic metric space. Then X and $\mathrm{Con}(\partial X)$ are roughly similar.*

The proof will actually show more. If d is a metric in the canonical gauge on ∂X for which there exists a point $w \in X$, and a constant $a \geqslant 1$ such that

$$a^{-1} e^{-(z|z')_w} \leqslant d(z, z') \leqslant a e^{-(z|z')_w}, \quad z, z' \in \partial X,$$

then X and $\mathrm{Con}((\partial X, d))$ are roughly isometric.

Proof. Assume X is δ-hyperbolic and k-visual with respect to $o \in X$. If we take different metrics d_1, d_2 in the canonical gauge on ∂X, then the spaces $(\partial X, d_1)$ and $(\partial X, d_2)$ are snowflake equivalent. Theorem 7.4 implies that the spaces $\mathrm{Con}((\partial X, d_1))$ and $\mathrm{Con}((\partial X, d_2))$ are roughly similar. Therefore, the assertion does not depend on which metric in the canonical gauge on ∂X we choose. Fix such a metric d on ∂X. Then we may assume that

$$a^{-1} \exp\big(-\epsilon(z|z')\big) \leqslant d(z, z') \leqslant a \exp\big(-\epsilon(z|z')\big), \quad \forall z, z' \in \partial X. \quad (8.1)$$

Making k larger if necessary, by Proposition 5.6 and Proposition 5.2.(2) for each $z \in \partial Z$ we can choose a k-roughly geodesic ray $\gamma_z : [0, \infty) \to X$ with $\gamma(o) = 0$ and $\lim \gamma_z = z$.

Let $D = D((\partial X, d))$. For $(z, h) \in \mathrm{Con}(\partial X)$, we define $f(z, h) = \gamma_z(\epsilon^{-1} \log(D/h))$. We show that $f : \mathrm{Con}(\partial X) \to X$ is a rough similarity.

First, note that the image of f is cobounded in X. For each $x \in X$ lies on some k-roughly geodesic ray emanating from o. If $z = \lim \gamma$, then from Lemma 5.1 it follows that $|\gamma(t) - \gamma_z(t)| = O_{\delta,k}(1)$ for $t \geqslant 0$. So x has distance at most $O_{\delta,k}(1)$ to γ_z which is contained in the image of f.

Let (z, h) and (z', h') be two arbitrary points in $\mathrm{Con}(\partial X)$. Observe that $(z|f(z, h)) = \epsilon^{-1} \log(D/h) + O_{\delta,k}(1)$, and $(z'|f(z', h')) = \epsilon^{-1} \log(D/h') + O_{\delta,k}(1)$. Therefore, from Lemma 5.1 we get

$$\begin{aligned}
\epsilon &\big| f(z, h) - f(z', h') \big| \\
&= \log(D/h) + \log(D/h') - 2\big(\epsilon(z|z') \wedge \log(D/h) \wedge \log(D/h')\big) + O_{\delta,k,\epsilon}(1) \\
&= -\log h - \log h' + 2\big(-\epsilon(z|z') \vee \log h \vee \log h'\big) + O_{\delta,k,\epsilon,D}(1) \\
&= 2 \log\left(\frac{e^{-\epsilon(z|z')} \vee h \vee h'}{\sqrt{hh'}}\right) + O_{\delta,k,\epsilon,D}(1).
\end{aligned}$$

Now apply (8.1) and the inequalities $(u + v)/2 \leqslant u \vee v \leqslant u + v$ to obtain

$$\epsilon \big| f(z, h) - f(z', h') \big| = 2 \log \left(\frac{d(z, z') + h \vee h'}{\sqrt{hh'}} \right) + O_{\delta, k, \epsilon, D, a}(1)$$

$$= \rho\big((z, h), (z', h')\big) + O_{\delta, k, \epsilon, D, a}(1). \qquad (8.2)$$

Since the image of f is cobounded in X, it follows from (8.2) that f is a rough quasi-similarity. Note that if $\epsilon = 1$, then f is a rough isometry.   □

The last two statements may be summarized as follows. If Z is a complete bounded metric space, then

$$Z \overset{\text{sf}}{=} \partial(\mathrm{Con}\, Z)\,,$$

where $\overset{\text{sf}}{=}$ means snowflake equivalence of the spaces. Note that for a space to be snowflake equivalent to the boundary of a Gromov hyperbolic space it is necessary that the space is complete and bounded by Proposition 6.2.

On the other hand, if X is a visual Gromov hyperbolic space, then

$$X \simeq \mathrm{Con}(\partial X)\,,$$

where $\simeq$ means rough similarity of the spaces. It is not hard to see that the property of being visual and Gromov hyperbolicity are preserved under rough similarities. So the assumptions on X are necessary by Theorem 7.2.

These two statements show that in some sense $X \mapsto \partial X$ and $Z \mapsto \mathrm{Con}(Z)$ are inverse to each other.

9   Growth and Assouad Dimension

In this section we prove Theorem 9.2 which will be used in sections 10 and 11.

DEFINITION (cf. [A]).   Let X be a metric space. For $\alpha, \beta > 0$, let $S(\alpha, \beta)$ be the maximal cardinality of a set $V \subset X$ such that $\alpha \leqslant |x - y| \leqslant \beta$ for all $x, y \in V$, $x \neq y$. Define t to be the infimum of all numbers $s \geqslant 0$ such that for some constant $K \geqslant 0$ the inequality

$$S(\alpha, \beta) \leqslant K(\beta/\alpha)^s$$

holds for all $0 < \alpha \leqslant \beta$. (It is understood that $t = \infty$, if no such s exists.) Then $\dim_A(X) = t$ is the *Assouad dimension* of X.

See [L] for a further discussion of the Assouad dimension. We will need the following theorem (cf. [A]).

Theorem 9.1 (Assouad's Theorem).   *Let (Z, d) be a metric space with finite Assouad dimension, and let $p \in (0, 1)$. Then there is some integer n such that the metric space (Z, d^p) admits a bilipschitz embedding into $\mathbb{R}^n$.*

Here, d^p is the metric $d^p(x,y) = d(x,y)^p$, $x,y \in Z$.

A metric space X is called *doubling*, if for all r, R with $R > r > 0$ there exists $N \in \mathbb{N}$ only depending on R/r such that every open ball of radius R in X can be covered by N open balls of radius r. It can be shown that a metric space has finite Assouad dimension if and only if it is doubling.

The next theorem gives a sufficient condition for a Gromov hyperbolic geodesic metric space X to have a boundary with finite Assouad dimension. To state the theorem we make the following definition.

DEFINITION. A metric space X has *bounded growth at some scale*, if there are constants r, R with $R > r > 0$, and $N \in \mathbb{N}$ such that every open ball of radius R in X can be covered by N open balls of radius r.

Theorem 9.2. *Let X be a Gromov hyperbolic geodesic metric space with bounded growth at some scale. Then the Assouad dimension of ∂X is finite.*

As we remarked above, the assertion is equivalent to ∂X being doubling. So a weak doubling condition on X implies that ∂X is doubling. Note that the properties of a metric space to have finite Assouad dimension and to be doubling are preserved by quasisymmetric maps. In particular, the statement about ∂X in the theorem is independent of which metric in the canonical gauge we choose.

Examples for Gromov hyperbolic geodesic metric spaces having bounded growth at some scale are all complete simply-connected Riemannian n-manifolds X with sectional curvature κ satisfying $-b^2 \leqslant \kappa \leqslant -a^2 < 0$. In this case open balls of radius 1 in X are bilipschitzly equivalent to the unit ball in $\mathbb{R}^n$ with bilipschitz constant only depending on a and b. This statement follows from Topogonov's comparison theorem, and implies that X has bounded growth at some scale. In particular, complex hyperbolic spaces are Gromov hyperbolic geodesic metric spaces with bounded growth at some scale.

Note that any bounded degree graph has bounded growth at some scale. Hence, every finitely generated hyperbolic group has bounded growth at some scale.

The proof of 9.2 is similar to the proof of Proposition 11 from [GH, Ch. 7].

Proof. Assume X is δ-hyperbolic, and fix some metric $d_{\partial X} = d_{o,\epsilon}$ in the canonical gauge on ∂X. Rescaling the metric on X by the factor ϵ if necessary, we may assume $\epsilon = 1$. Then

$$(z|z') = -\log\big(d_{\partial X}(z,z')\big) + O_\delta(1), \qquad \forall z, z' \in \partial X. \qquad (9.1)$$

Denote by $B(q, s)$ the open ball in X with center $q \in X$ and radius $s > 0$. By assumption, there exist constants $R > r > 0$, $N \in \mathbb{N}$ such that every open ball of radius R in X can be covered by N open balls of radius r. Suppose that $B(p, 2R - r)$ is some ball of radius $2R - r$ and center $p \in X$. Then $B(p, R)$ can be covered by N open balls of radius r with centers $p_1, \ldots, p_N \in X$, say. Since X is a geodesic metric space, it follows that $B(p, 2R - r)$ is covered by the balls $B(p_1, R), \ldots, B(p_N, R)$. Therefore, $B(p, 2R - r)$ can be covered by N^2 open balls of radius r. By induction it follows that any open ball of radius $mR - (m-1)r$, $m \in \mathbb{N}$, can be covered by N^m open balls of radius r.

Given any $z \in \partial X$, for every $t > 0$ there is an $x \in X$ such that $|x| = t$ and $(x|z) - t = O_\delta(1)$. To verify this, take $w \in X$ with $(w|z) \geqslant t$ and let x be the point satisfying $|x| = t$ on the geodesic segment $[o, w]$.

Fix some α, β with $0 < \alpha \leqslant \beta \leqslant 1$. Suppose $z_1, z_2, \ldots, z_n$ are points in ∂X such that

$$\alpha \leqslant d_{\partial X}(z_i, z_j) \leqslant \beta, \quad i \neq j. \tag{9.2}$$

Then (9.1) implies

$$-\log \beta - C(\delta) \leqslant (z_i|z_j) \leqslant -\log \alpha + C(\delta), \qquad i \neq j. \tag{9.3}$$

Let $x_1, \ldots, x_n$ be points satisfying $|x_j| = -\log \beta$, $\big| -\log \beta - (x_j|z_j) \big| \leqslant C(\delta)$, and let $y_1, \ldots, y_n$ be points satisfying $|y_j| = -\log \alpha$, $\big| -\log \beta - (y_j|z_j) \big| \leqslant C(\delta)$. Then it is easy to see from Lemma 5.1 and (9.3) that

$$
\begin{aligned}
|x_i - x_j| &= O_\delta(1), \\
|y_i - y_j| &= 2(-\log \alpha + \log \beta) + O_\delta(1), \\
|x_i - y_i| &= \log \beta - \log \alpha + O_\delta(1),
\end{aligned}
\tag{9.4}
$$

for all $i \neq j$. It follows that the open ball of radius

$$R^* = \log(\beta/\alpha) + C(\delta)$$

centered at x_1 contains all the points y_i.

Let m be the least integer such that $m(R - r) \geqslant R^*$. We have seen that an open ball of radius $m(R - r)$ can be covered by N^m balls of radius r. So the set $\{y_1, \ldots, y_n\}$ can be covered by N^m balls of radius r. We now assume that $\log(\beta/\alpha)$ is much larger than r. It then follows from the middle equation in (9.4) that $|y_i - y_j| > 2r$ for $i \neq j$. Consequently, any ball of radius r can contain at most one of the points y_j. This gives

$$n \leqslant N^m \leqslant N^{1 + R^*/(R-r)} \leqslant C(R, r, N, \delta)(\beta/\alpha)^{C(R, r, N, \delta)}.$$

By changing the constants if necessary, the same is true even without the assumption that $\log(\beta/\alpha)$ is much larger than r. Since the diameter of X

is finite, the assumption $\beta \leqslant 1$ is also inconsequential. We conclude that the Assouad dimension of ∂X is finite. $\square$

10   Embeddings into Real Hyperbolic Spaces

The definition of the metric ρ on $\mathrm{Con}(X)$ was motivated by the upper half-space model of real hyperbolic n-space $\mathbb{H}^n$. In this model

$$\mathbb{H}^n = \mathbb{R}^n_+ = \big\{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_n > 0\big\}$$

is equipped with the Riemannian metric given by the length element

$$ds^2 = \tfrac{1}{x_n^2}(dx_1^2 + \cdots + dx_n^2).$$

The space $\mathbb{H}^n$ is a Riemannian n-manifold with constant sectional curvature -1.

We identify $\mathbb{R}^{n-1}$ with the hyperplane in $\mathbb{R}^n$ given by the equation $x_n = 0$. For $x \in \mathbb{R}^n$ we write $x = (z, h)$, where $z \in \mathbb{R}^{n-1}$, $h \in \mathbb{R}$. Then $\mathbb{H}^n = \{(z, h) : z \in \mathbb{R}^{n-1}, h > 0\}$.

It is not hard to find explicit expressions for the hyperbolic distance $d_h((z, h), (z', h'))$ of two points $(z, h), (z', h') \in \mathbb{H}^n$. For our purpose it is sufficient to know that for some universal constant $C > 0$

$$\left| d_h\big((z, h), (z', h')\big) - 2\log\left(\frac{|z - z'| + h \vee h'}{\sqrt{hh'}}\right)\right| \leqslant C. \quad (10.1)$$

Here $|z - z'|$ is the euclidean distance of z and z'.

The spaces $\mathbb{H}^n$, $n \in \mathbb{N}$, are geodesic, and from (10.1) and the first part of the proof of Theorem 7.2 it follows that they are C-hyperbolic.

The Gromov boundary $\partial\mathbb{H}^n$ can be identified with $\mathbb{R}^{n-1} \cup \{\infty\}$. If $M \subset \mathbb{R}^n_+$ is bounded with respect to the euclidean metric, then (M, d_h) is C-hyperbolic, and ∂M can be identified with the intersection of the euclidean closure of M and $\mathbb{R}^{n-1}$. Moreover, if $w \in \mathbb{H}^n$, then

$$(z|z')_w = -\log|z - z'| + O_{M,w}(1), \quad \forall z, z' \in \partial M \subset \mathbb{R}^{n-1}. \quad (10.2)$$

In particular, the euclidean metric is bilipschitz to a metric in the canonical gauge of ∂M. These statements can be verified using (10.1), and considerations as in the proof of Theorem 8.1.

If $A \subset \mathbb{H}^n \cup \partial\mathbb{H}^n$, let $\mathrm{hull}(A) \subset \mathbb{H}^n$ be the intersection of all closed half-spaces $H \subset \mathbb{H}^n$ such that $A \subset H \cup \partial H$. Recall that in the upper half-space model of $\mathbb{H}^n$, half-spaces in $\mathbb{H}^n$ are determined by hyperplanes in $\mathbb{R}^n$ perpendicular to $\mathbb{R}^{n-1}$ or spheres with center in $\mathbb{R}^{n-1}$. The set $\mathrm{hull}(A)$ is the smallest set $M \subset \mathbb{H}^n$ that is hyperbolically convex and closed, and

satisfies $A \subset M \cup \partial M$. Here we have to assume that A consists of more than one point if $A \subset \partial \mathbb{H}^n$, for otherwise $\text{hull}(A) = \emptyset$.

We need the following facts about $\text{hull}(A)$.

PROPOSITION 10.1.   (1) *Let $A \subset \mathbb{H}^n$. For each $p \in \text{hull}(A)$, there exists a point $q \in \mathbb{H}^n$ that lies on a geodesic segment with endpoints in A and satisfies $d_h(p, q) \leqslant O(1)$.*

(2) *Let $A \subset \partial \mathbb{H}^n$ be a set with more than one point, and fix $o \in \text{hull}(A)$. For each $p \in \text{hull}(A)$, there exists a point $q \in \mathbb{H}^n$ that lies on a geodesic ray from o to some point in A and satisfies $d_h(p, q) \leqslant O(1)$.*

(3) *Let $Z \subset \mathbb{R}^{n-1} \subset \partial \mathbb{H}^n$ be a compact set containing more than one point. Then $\text{hull}(Z) \cong \text{Con}(Z)$.*

Remember the notation $X \sim Y$, if two metric spaces X, Y are roughly quasi-isometric, $X \simeq Y$ if they are roughly similar, and $X \cong Y$, if there are roughly isometric.

It is worthwhile to note that the constant $O(1)$ in the statement is absolute. In particular, it does not depend on the dimension n.

Proof. Statement (1) and (2) can easily be proved, for example, by using the Euclidean unit ball in $\mathbb{R}^n$ as the Klein model for $\mathbb{H}^n$ (where the hyperbolic geodesics are line segments) with p corresponding to 0. We leave the details to the reader.

(3) Since Z is compact, $\partial(\text{hull}(Z)) = Z$. To see this note that $Z \subset \partial(\text{hull}(Z))$. On the other hand, assume $z \in \mathbb{R}^{n-1} \backslash Z \subset \partial \mathbb{H}^n$. Since Z is compact, there exists a small closed euclidean ball B centered at z which is disjoint from Z. Then $H = B \cap \mathbb{R}^n_+$ is a closed half-space in $\mathbb{H}^n$ disjoint from $\text{hull}(Z)$. Hence, $z \notin \partial(\text{hull}(Z))$.

By the second part of (1), the union of the geodesic rays in $\text{hull}(Z)$ emanating from some fixed basepoint $o \in \text{hull}(Z)$ is cobounded. Therefore, $\text{hull}(Z)$ is visual.

Since $\text{hull}(Z)$ is also Gromov hyperbolic, Theorem 8.2 shows that $\text{hull}(Z) \simeq \text{Con}(Z)$. To get the stronger statement $\text{hull}(Z) \cong \text{Con}(Z)$ note that the euclidean metric on $Z = \partial(\text{hull}(Z))$ is bilipschitz to a metric in the canonical gauge on $\partial(\text{hull}(Z))$ by formula (10.2). The assertion now follows from the remark after Theorem 8.2. $\square$

Theorem 10.2.   *Let X be a Gromov hyperbolic geodesic metric space with bounded growth at some scale. Then there exists an integer n such that X is roughly similar to a convex subset of hyperbolic n-space $\mathbb{H}^n$.*

Note that the theorem has an easy converse. If X is geodesic and

roughly similar to a subset of $\mathbb{H}^n$, then X is Gromov hyperbolic and has bounded growth at some scale.

Let d be the metric on X. The theorem says that for some $c > 0$ the metric space (X, cd) is roughly isometric to a convex subset of $\mathbb{H}^n$. The proof shows that there is a $C = C(\delta(X)) > 0$ such that c may be chosen as any number in $(0, C)$. However, one cannot always take $c = 1$, as demonstrated by Proposition 10.3 below.

Proof. First suppose that the union of the geodesic rays starting from some basepoint $o \in X$ is cobounded in X, and that ∂X contains more than one point.

Then X is visual, and so $X \simeq \mathrm{Con}((\partial X, d))$ by Theorem 8.2. Here, d is some fixed metric in the canonical gauge on ∂X. By Theorem 7.4 applied to snowflake maps, $\mathrm{Con}((\partial X, d)) \simeq \mathrm{Con}((\partial X, d^{1/2}))$. Theorem 9.2 shows that ∂X has finite Assouad dimension. By Theorem 9.1, $(\partial X, d^{1/2})$ admits a bilipschitz embedding into $\mathbb{R}^{n-1}$ for sufficiently large n. Let Z be the bilipschitz image of $(\partial X, d^{1/2})$ in $\mathbb{R}^{n-1}$. By Theorem 7.4 for bilipschitz maps, $\mathrm{Con}((\partial X, d^{1/2})) \cong \mathrm{Con}(Z)$. Since ∂X is complete, bounded, and contains more than one point, Z also has these properties. In particular, Z is compact. Proposition 10.1 shows that $\mathrm{Con}(Z) \cong \mathrm{hull}(Z)$. Therefore, $X \simeq \mathrm{hull}(Z) \subset \mathbb{H}^n$, proving the assertion in this case.

Now drop the additional assumptions on X. Since X has bounded growth at some scale, there exist constants $R > r > 0$, $N \in \mathbb{N}$ such that every open ball of radius R in X can be covered by N open balls of radius r. Let X_0 be a maximal set of points in X with the property that the distance between any two points in X_0 is at least $5R$.

At each point $x_0 \in X$, we glue an isometric copy of the ray $[0, \infty)$ to X by identifying x_0 with 0, the initial point of the ray. The new space $\widehat{X}$ carries a unique metric which agrees with the metric on X, the metrics on the rays glued to X, and makes $\widehat{X}$ a geodesic space. If a geodesic segment connects two points lying on different rays glued to X, then this segment contains the initial points of the rays. Using this and the thin triangle definition of hyperbolicity, it is not hard to check that $\widehat{X}$ is Gromov hyperbolic.

If we assign to each point $x_0 \in X_0$ the limit point in $\partial \widehat{X}$ of the ray glued to x_0, then we get an injective embedding of X_0 into $\partial \widehat{X}$. From this we see that we may assume that $\partial \widehat{X}$ contains more than one point. For otherwise, X_0 is a one point set. Then X has bounded diameter. The theorem is true in this case, X being roughly similar to a point.

Fix some basepoint $o \in X$. For each $x_0 \in X_0$, the union of a geodesic segment $[o, x_0]$, and the ray glued to X at x_0 is a geodesic ray emanating from o. By definition of X_0 this implies that the union of the geodesic rays from o is $5R$-cobounded.

To verify that $\widehat{X}$ has bounded growth at some scale, observe that a ball of radius R in $\widehat{X}$ can intersect at most one component of $\widehat{X} - X$. Consequently, any open ball of radius R in $\widehat{X}$ can be covered by $N + \lceil R/r \rceil$ open balls of radius r, where $\lceil R/r \rceil$ denotes the least integer greater than R/r.

It follows that $\widehat{X}$ satisfies the additional assumptions of the special case treated above. So we know that $\widehat{X} \simeq \widehat{W}$ for some convex set $\widehat{W} \subset \mathbb{H}^n$, where n is sufficiently large. Let W denote the image of X under a rough similarity $f : \widehat{X} \to \widehat{W}$. Then $X \simeq W$.

We claim that $W \cong \mathrm{hull}(W)$. To see this let $w, w' \in W$. There exist points $x, x' \in X$ such that $w = f(x)$ and $w' = f(x')$. Since X is geodesic, there exists a geodesic segment $[x, x']$ joining x and x'. The image $f([x, x'])$ is a roughly quasi-isometric path joining w and w'. By geodesic stability in $\mathbb{H}^n$, $d_{\mathrm{Haus}}(f([x, x']), [w, w']) \leqslant c$, where c depends on the parameters of f only.

This shows that $W = f(X)$ is cobounded in the set consisting of the union of all geodesic segments with endpoints in W. Since the latter set is cobounded in $\mathrm{hull}(W)$ by Proposition 10.1, the set W is also cobounded in $\mathrm{hull}(W)$. Therefore, the inclusion map $i : W \to \mathrm{hull}(W)$ is a rough isometry. We conclude $X \simeq \mathrm{hull}(W) \subset \mathbb{H}^n$, completing the proof of the theorem. $\qquad\square$

The following theorem shows that in 10.2 "roughly similar" cannot be replaced by "roughly isometric".

PROPOSITION 10.3.   *Let d_h denote the standard metric on $\mathbb{H}^2$, let $c > 1$ and $n \in \mathbb{N}$. Then the metric space $X = (\mathbb{H}^2, c\, d_h)$ is not roughly isometric to a subset of $\mathbb{H}^n$.*

Proof. Recall that $\partial \mathbb{H}^n$ the $(n-1)$-dimensional unit sphere is S^{n-1}, in the following sense. For every $o \in \mathbb{H}^n$ the metric $d_{\partial \mathbb{H}^n, o, 1}$ is bilipschitz to $d_{S^{n-1}}$, the spherical metric on S^{n-1}. We identify $\partial \mathbb{H}^n$ with S^{n-1}.

Suppose that X is roughly isometric to a subset W of $\mathbb{H}^n$. Then it must be roughly isometric to $\mathrm{hull}(W)$. Fix points $o \in X$, $o' \in W$, and fix a small $\epsilon > 0$. It follows that the metric $d_{\partial \mathbb{H}^2, o, c\epsilon}$ is bilipschitz to the metric $d_{\partial \mathrm{hull}(W), o', \epsilon}$, which is just the restriction of $d_{\partial \mathbb{H}^n, o', \epsilon}$ to ∂W. It therefore follows that $(d_{S^1})^c$ is bilipschitz to the restriction of $d_{S^{n-1}}$ to ∂W, in the

sense that there is a homeomorphism $g : S^1 \to \partial W \subset S^{n-1}$ and a constant C such that

$$C^{-1} d_{S^{n-1}}\big(g(x), g(y)\big) \leqslant d_{S^1}(x, y)^c \leqslant C d_{S^{n-1}}\big(g(x), g(y)\big),$$

for every $x, y \in S^1$. However, for every $x, y \in S^1$ and every $s > 0$ there is a sequence $x = x_0, x_1, \ldots, x_k = y$ such that $\sum_{j=1}^{k} d_{S^1}(x_{j-1}, x_j)^c < s$. It follows that $d_{S^1}(x, y) = 0$, a contradiction. $\square$

It is interesting to ask what can remain of the conclusion of Theorem 10.2 if the assumption that X has finite growth at some scale is removed.

For any metric space Z let

$$\mathrm{Con}_h(Z) = Z \times (0, \infty),$$

and define a metric d_h on $\mathrm{Con}_h(Z)$ as follows. Given two points $p = (z, h), p' = (z', h') \in \mathrm{Con}_h(Z)$, let $d_h(p, p')$ be the distance between the two points in the hyperbolic plane $\mathbb{H}^2$ whose coordinates in the upper half-plane model are $(0, h)$ and $(|z - z'|, h')$. Note that $\mathrm{Con}_h(\mathbb{R}^{n-1}) = \mathbb{H}^n$. The spaces $\mathrm{Con}_h(Z)$ are δ-hyperbolic for some univeral $\delta \geqslant 0$. This is seen by checking the four point condition (3.3) using (10.1), and the argument of the first part of the proof of Theorem 7.2.

If Y is any bounded subset of Z, then (10.1) implies that the set inclusion $i : \mathrm{Con}(Y) \to \mathrm{Con}_h(Z)$ is a roughly isometric embedding of $\mathrm{Con}(Y)$ into $\mathrm{Con}_h(Z)$.

On first thought it seems tempting to expect that perhaps any Gromov hyperbolic space admits a rough similarity into $\mathrm{Con}_h(H)$, where H is a sufficiently large Hilbert space. This seems to be natural, since $\mathrm{Con}_h(H)$ of a Hilbert space H is the infinite dimensional analogue of $\mathbb{H}^n$. However, it follows from results of Enflo [E] (see also [Ra]) that there is a bounded metric spaces (Z, d) such that none of the spaces (Z, d^p), $p \in (0, 1]$, has a bilipschitz embedding into any Hilbert space. This can be used to show that $\mathrm{Con}(Z)$ does not admit a rough similarity into $\mathrm{Con}_h(H)$ for any Hilbert space H.

On the other hand, the following is easy.

Theorem 10.4.   *For every Gromov hyperbolic space X there is an l_∞-space $l_\infty(A)$ and a rough similarity of X into $\mathrm{Con}_h(l_\infty(A))$.*

Recall that if A is any set, then $l_\infty(A)$ is the space of all bounded real-valued functions on A equipped with the metric coming from the supremum norm.

Proof. Using Theorem 4.1, we assume with no loss of generality that X is a geodesic metric space. Choose a basepoint $o \in X$. To every point $x \in X$ glue a new copy of the ray $[0, \infty)$, with the initial point of the ray identified with x, and call the resulting space $\widehat{X}$. Endow $\widehat{X}$ with the obvious metric, and note that $\widehat{X}$ is Gromov hyperbolic. Then $\widehat{X}$ has the property that for every point $x \in \widehat{X}$ there is a geodesic ray from o passing through x (cf. the proof of Theorem 10.2). Let $Z = \partial \widehat{X}$. Then $\widehat{X} \simeq \text{Con}(Z)$ by Theorem 8.2. Fix a metric d in the canonical gauge on Z. Then the map $f : (Z, d) \rightarrow l_\infty(Z)$, $z \mapsto d(z, \cdot)$, is an isometric embedding of (Z, d) into $l_\infty(Z)$. Thus $\text{Con}(Z) \cong \text{Con}(f(Z))$. Since the inclusion $\text{Con}(f(Z)) \subset \text{Con}_h(l_\infty(Z))$ is a rough isometric embedding by the remark following the definition of Con_h, $\widehat{X}$ admits an embedding into $\text{Con}_h(l_\infty(Z))$ by a rough similarity. The same is then true for X. $\square$

11   Some Two-dimensional Applications

In this section a characterization of the hyperbolic plane among geodesic metric spaces up to rough quasi-isometry is given. It would be interesting to have similar statements for higher dimensional hyperbolic spaces as well. We also give a sufficient condition for a Gromov hyperbolic planar graph to be roughly quasi-isometric to a convex subset of $\mathbb{H}^2$.

DEFINITION.   A metric space Z is called a λ-*quasi-circle* for $\lambda \geqslant 1$, if it is homeomorphic to the circle $S^1 = \{x \in \mathbb{C} : |x| = 1\}$, and has the property that for any pair of distinct points $x, y \in Z$, the diameter of at least one of the two components of $Z - \{x, y\}$ is at most $\lambda|x - y|$.

We need the following result (cf. [TuV1, Thm. 4.9]).

Theorem 11.1. *A metric space Z is power quasisymmetric to S^1 if and only if it is a quasi-circle and doubling.*

The theorem was not exactly stated like this in [TuV1], but our statement is equivalent. To see this note that a metric space is doubling if and only if it has finite Assouad dimension, and this is equivalent to the condition that the space is homogeneously totally bounded as defined in [TuV1] (cf. Def. 2.7 and Rem. 3.20). Moreover, every quasisymmetry between connected metric spaces is a power quasisymmetry (cf. [TuV1, Cor. 3.12]).

There are quasi-circles which are not doubling. The unit circle S^1 with the metric $d(x, y) = 1/\log(100/|x - y|)$, $x \neq y$, is an example.

Theorem 11.2. *A geodesic metric space X is roughly quasi-isometric to*

the hyperbolic plane $\mathbb{H}^2$ *if and only if* X *is Gromov hyperbolic, visual, has bounded growth at some scale, and* ∂X *is a quasi-circle.*

The property of a metric space being a quasi-circle is preserved by quasisymmetric maps. In particular, for the condition that ∂X is a quasi-circle it does not matter which metric in the canonical gauge on ∂X we choose.

There are Gromov hyperbolic geodesic metric spaces which are visual, have bounded growth at some scale, and for which ∂X is a topological circle, but which are not roughly quasi-isometric to $\mathbb{H}^2$. For example, one can take X to be the hyperbolic convex hull of a topological circle Z in $\partial \mathbb{H}^3$ which is not a quasi-circle. If X and $\mathbb{H}^2$ were roughly quasi-isometric, then $Z = \partial X$ and $\partial \mathbb{H}^2$ would be power quasisymmetric. This is impossible, since $\partial \mathbb{H}^2$ is a quasi-circle (cf. the proof below).

Proof. First note that $\mathbb{H}^2$ is a Gromov hyperbolic geodesic metric space which is visual and has bounded growth at some scale. If we use the unit disc model of the hyperbolic plane, then $\partial \mathbb{H}^2$ can be identified with the unit circle S^1, and it can be shown that the euclidean metric on S^1 is bilipschitz to a metric in the canonical gauge. Therefore, $\partial \mathbb{H}^2$ is a quasi-circle.

Conversely, if a metric space has all the properties stated in the theorem, then $X \simeq \mathrm{Con}(\partial X)$ by Theorem 8.2. By Theorem 9.2 the Assouad dimension of ∂X is finite, and so ∂X is doubling. Hence ∂X is power quasisymmetric to S^1 by Theorem 11.1. By the first part of the proof, we can apply these considerations to $\mathbb{H}^2$. In particular, $\mathbb{H}^2 \simeq \mathrm{Con}(\partial \mathbb{H}^2)$. Moreover, both $\partial \mathbb{H}^2$ and ∂X are power quasisymmetric to S^1, and so there are power quasisymmetric to each other. This implies $\mathrm{Con}(\partial X) \sim \mathrm{Con}(\partial \mathbb{H}^2)$ by Theorem 7.4. We conclude that $X \sim \mathbb{H}^2$.   □

Following is an application of Theorem 11.2 to planar graphs. We consider only connected graphs G and think of G as being equipped with a path metric in the usual way, where the edges are intervals of length 1.

Theorem 11.3.   *Let* G *be a Gromov hyperbolic planar graph with bounded vertex degree. Suppose that the union of all geodesics is cobounded in* G, *and that every compact subset of the plane contains only finitely many vertices of* G. *Then* G *is roughly quasi-isometric to a convex subset of* $\mathbb{H}^2$.

The following lemma is probably known, but we have not found a reference.

LEMMA 11.4.   *Suppose that a geodesic metric space* Z *is the union,* $Z = X \cup Y$, *of two closed subsets* $X, Y \subset Z$. *Assume that* $X \cap Y$ *is*

geodesic, and both X and Y are δ-hyperbolic. Then Z is δ'-hyperbolic, where δ' only depends on δ.

A proof can be based on Rips' thin triangles definition. The details are left to the reader.

We shall also need a recent theorem of P. Bowers.

Theorem 11.5 [B]. *Let G be the 1-skeleton of a triangulation of the plane, which is Gromov hyperbolic. Then the Gromov boundary ∂G is either a single point or a topological circle.*

In the latter case, suppose that γ is a geodesic in G, and let G_1, G_2 be the two halves of G "cut" by γ. Then G_1 and G_2 are Gromov hyperbolic and the closed arcs on the circle ∂G determined by the endpoints of γ are the Gromov boundaries of G_1 and G_2.

Proof of 11.3. Assume that G is δ-hyperbolic.

Our first goal is to obtain a better understanding of the planar embedding of G. A *face* is a component of $\mathbb{R}^2 - G$. Let F be a bounded face. Then there are finitely many vertices of G in ∂F. We shall now show that the number of vertices in ∂F is bounded independently of F.

Given any simple closed path α in G, let $D(\alpha)$ denote the closure of the bounded component of $\mathbb{R}^2 - \alpha$. Let n be the least number of vertices in a simple closed path α in G satisfying $D(\alpha) \supset F$. Let α be a simple closed path in G with n vertices such that $D(\alpha) \supset F$ and with $D(\alpha)$ as small as possible. That is, there is no $\alpha' \neq \alpha$ with n vertices and $F \subset D(\alpha') \subset D(\alpha)$. If we choose three almost equally spaced vertices in α, then α becomes a geodesic triangle. Therefore, from the Rips thin triangles condition it follows that there is an upper bound of $O_\delta(1)$ on n, the number of vertices in α.

No geodesic γ in G intersects the interior of $D(\alpha)$, because otherwise we could contradict the minimality of the length of α or of $D(\alpha)$ by replacing an arc of α by an arc of $\gamma \cap D(\alpha)$. Since the union of all geodesics is cobounded G, it follows that there is an upper bound on the distance from any vertex on ∂F to α. Since the degrees of the vertices in G are bounded, and the length n, of α is bounded, it follows that the number of vertices in ∂F is bounded. Consequently, we assume, with no loss of generality, that each bounded face is a triangle (that is, has three vertices on its boundary), because triangulating the bounded faces gives a planar graph roughly isometric to G.

We now deal with the unbounded faces, and prove that with no loss of generality it may be assumed that the boundary of an unbounded face

is a geodesic. Let F be an unbounded face. Then there are infinitely many vertices in ∂F. Fix two vertices a, b in ∂F, and let β be a simple curve in the interior of F with endpoints a, b. Then there is a geodesic arc $\alpha_{a,b}$ joining a and b such that $D(\alpha_{a,b} \cup \beta)$ is minimal with respect to inclusion. It follows that no geodesic enters $D(\alpha_{a,b} \cup \beta)$. Now let a and b tend to infinity along ∂F, "in opposite directions". Since the union of the geodesics is cobounded in G, there is an upper bound on the lower limit of the combinatorial distance of an arbitrary vertex $v \in \partial F$ from $\alpha_{a,b}$. One consequence is that we may extract a limiting geodesic α_F. It also follows that there are no geodesics that meet the component of $\mathbb{R}^2 - \alpha_F$ which contains F. Let G_F be the set of vertices and edges of G that are contained in the component of $\mathbb{R}^2 - \alpha_F$ that contains F, and let G' be the graph obtained from G by deleting G_F for every unbounded face F. Then G' is roughly isometric to G. Consequently, assume with no loss of generality that $G = G'$ and for every unbounded F, its boundary ∂F is a geodesic in G.

Now let Q be a bounded degree triangulation of the upper half-plane that is roughly quasi-isometric to a hyperbolic half-plane and such that ∂Q, the part of Q on the real axis, is a geodesic. For each unbounded face F of G, glue a graph Q_F combinatorially isomorphic to Q, with ∂Q_F identified with the geodesic ∂F, and with orientations agreeing. Let $\widehat{G}$ be the resulting graph.

Suppose that x_1, x_2, x_3, x_4 are four vertices in $\widehat{G}$. Then there are at most four unbounded faces F such that $\{x_1, x_2, x_3, x_4\}$ is contained in G union with the corresponding Q_F's. Consequently, it follows by applying Lemma 11.4 three times that the distances determined by these four points satisfy (3.3) with some constant δ' independent of the choice of the points x_1, x_2, x_3, x_4. Therefore, $\widehat{G}$ is Gromov hyperbolic. Moreover, it is (combinatorially isomorphic to) a triangulation of the plane where the union of the geodesics is cobounded, and G is isometric with a subset of $\widehat{G}$.

From Theorem 11.5 it follows that the Gromov boundary $\partial\widehat{G}$ is a topological circle. We will show that it is a quasi-circle.

Fix a basepoint $o \in \partial\widehat{G}$. If $x, y \in \partial\widehat{G}$, $x \neq y$, there exists a geodesic γ_{xy} in $\widehat{G}$ joining x and y, i.e., one end of γ_{xy} converges to x and one to y. It can be shown that

$$(x|y) = \operatorname{dist}(o, \gamma_{xy}) + O_\delta(1).$$

It follows that if d is some fixed metric in the canonical gauge on $\partial\widehat{G}$, then

there exists $\epsilon > 0$ such that

$$d(x,y) \asymp e^{-\epsilon \operatorname{dist}(o,\gamma_{xy})} \,. \tag{11.1}$$

Here $\asymp$ means equality up to multiplicative constants independent of x and y.

Let z_1, z_2 be two distinct points in $\partial \widehat{G}$, let $\alpha_1, \alpha_2 \subset \partial \widehat{G}$ be the two closed arcs on $\partial \widehat{G}$ determined by z_1, z_2, and let $w_1 \in \alpha_1, w_2 \in \alpha_2$ be points chosen to maximize $d(z_1, w_1)$ and $d(z_2, w_2)$. Note that $z_1 \neq w_1$, $z_2 \neq w_2$. Let $\gamma \subset \widehat{G}$ be a geodesic joining z_1 and z_2. Then $\widehat{G} - \gamma$ has two components, one, A_1 say, that contains a geodesic ray with limit point w_1, and one, A_2 say, that contains a geodesic ray with limit point w_2.

Changing notation if necessary, we may assume that $o \in \gamma \cup A_1$. The geodesic $\beta \subset X$ joining w_2 and z_2 lies entirely in $\gamma \cup A_2$. In order to reach β from o, one must cross γ. Therefore, $\operatorname{dist}(o, \gamma) \leqslant \operatorname{dist}(o, \beta)$. By (11.1) this implies

$$d\text{-}\operatorname{diam}(\alpha_2) \leqslant 2d(z_2, w_2) \leqslant c_1 e^{-\epsilon \operatorname{dist}(o,\beta)}$$
$$\leqslant c_1 e^{-\epsilon \operatorname{dist}(o,\gamma)} \leqslant c_2 d(z_1, z_2) \,.$$

Here the constants $c_1, c_2 > 0$ can be chosen independently of the points z_1, z_2, w_1, w_2. This shows that $\partial \widehat{G}$ is a quasi-circle.

The graph $\widehat{G}$ is visual, as easily follows from the fact that the union of the geodesics is cobounded in $\widehat{G}$. Since $\widehat{G}$ has bounded degree, it has bounded growth at some scale, and $\partial \widehat{G}$ has finite Assouad dimension. Theorem 11.2 then tells us that $\widehat{G}$ is roughly quasi-isometric to $\mathbb{H}^2$. Consequently, G is quasi-isometric to a subset, say W of $\mathbb{H}^2$. As in the proof of Theorem 10.2, it can be shown that W is roughly isometric to $\operatorname{hull}(W)$. The theorem follows. $\qquad \square$

COROLLARY 11.6.   *Assume additionally in Theorem 11.3 that ∂G is connected. Then G is roughly quasi-isometric to the hyperbolic plane, or to a hyperbolic half-plane.*

The proof easily follows from Theorem 11.3, and is left to the reader.

References

[A]   P. ASSOUAD, Plongements lipschitziens dans $\mathbb{R}^n$, Bull. Soc. Math. France 111 (1983), 429–448.

[B]   P.L. BOWERS, Negatively curved graph and planar metrics with applications to type, Michigan Math. J. 45 (1998), 31–53.

[CDP] M. COORNAERT, T. DELZANT, A. PAPADOPOULOS, Geométrié et théorie des groupes. Les groupes hyperboliques de Gromov, Springer Lecture Notes in Mathematics 1441, Berlin, 1990.

[E]   P. ENFLO, On a problem of Smirnov, Ark. Mat. 8 (1969), 107–109.

[GH]   E. GHYS, P. DE LA HARPE, EDS., Sur les Groupes Hyperboliques d'après Mikhael Gromov, Progress in Mathematics 83, Birkhäuser, Boston, 1990.

[Gr]   M. GROMOV, Hyperbolic groups, in "Essays in Group Theory" (G. Gersten, ed.), Math. Sci. Res. Inst. Publ. Springer (1987), 75–263.

[L]   J. LUUKKAINEN, Assouad dimension, antifractal metrization, porous sets, and homogeneous measures, J. Korean Math. Soc. 35 (1998), 23–76.

[P]   F. PAULIN, Un groupe hyperbolique est déterminé par son bord, J. London Math. Soc. 54 (1996), 50–74.

[R]   J.G. RATCLIFFE, Foundations of Hyperbolic Manifolds, Graduate Texts in Mathematics 149, Springer, New York, 1994.

[Ra]   Y. RAYNAUD, Espaces de Banach superstables, distances stables et homéomorphismes uniformes, Israel J. Math. 44 (1983), 33–52.

[TV]   D.A. TROTSENKO, J. VÄISÄLÄ, Upper sets and quasisymmetric maps, preprint.

[TuV1] P. TUKIA, J. VÄISÄLÄ, Quasisymmetric embeddings of metric spaces, Ann. Acad. Sci. Fenn. Ser. A I Math. 5 (1980), 97–114.

[TuV2] P. TUKIA, J. VÄISÄLÄ, Quasiconformal extension from dimension n to $n+1$, Ann. Math. 115 (1982), 331–348.

[V]   J. VÄISÄLÄ, Quasisymmetric embeddings in Euclidean spaces, Trans. Amer. Math. Soc. 264 (1981), 191–204.

MARIO BONK, Institut für Analysis, Technische Universität Braunschweig, D-38106 Braunschweig, Germany
M.Bonk@tu-bs.de

ODED SCHRAMM, Mathematics Department, The Weizmann Institute, Rehovot 76100, Israel
schramm@wisdom.weizmann.ac.il

Correction to "Embeddings of Gromov hyperbolic spaces". The parts of the proof of Theorem 9.2 related to the relations (9.4) are somewhat garbled, but it is easy to correct the argument. The idea is to find for each of the points z_j suitable associated points x_j and y_j in X. The points $x_1, \ldots, x_n$ should be close together, the distance $|x_j - y_j|$ should be controlled, but we want the definite separation $|y_i - y_j| > 2r$ for $i \neq j$. As in the paper, one chooses x_j for $j = 1, \ldots, n$ such that $|x_j| = -\log \beta$ and $|-\log \beta - (x_j|z_j)| \leq C(\delta)$. Then $|x_i - x_j| = O_\delta(1)$ as follows from Lemma 5.1.

The point y_j is chosen such that $|y_j| = t$ and $|(y_j|z_j) - t| \leq C(\delta)$ for appropriate $t \geq 0$. If $t \geq -\log \alpha$, then again Lemma 5.1 shows that

$$|x_j - y_j| = t + \log \beta + O_\delta(1)$$

and

$$|y_i - y_j| \geq 2t + 2\log \alpha + O_\delta(1)$$

for $i \neq j$. If we pick $t = -2\log \alpha + \log \beta + C(\delta)$, then

$$|x_j - y_j| = 2(-\log \alpha + \log \beta) + O_\delta(1)$$

and

$$|y_j - y_i| > 2\log(\beta/\alpha).$$

Here we may assume that $\log(\beta/\alpha) \geq r$, because if this inequality is not true, then we replace α by a smaller quantity that is comparable to α up to a factor only depending on r. The rest of the argument is as in the paper.

Mario Bonk

1