δ-HYPERBOLIC SPACES

SIDDHARTHA GADGIL

ABSTRACT. These are notes for the Chennai TMGT conference on δ-hyperbolic spaces corresponding to chapter III.H in the book of Bridson and Haefliger.

When viewed from a distance, the lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$ looks like the plane. This is a simple instance of a very fruitful philosophy, introduced by Gromov, that viewed from a distance many groups can be regarded as geometric objects.

A very important class of such groups is those whose large-scale behaviour is similar to that of negatively curved spaces. The goal of this chapter is to lay the foundations for the study of these by making sense of the expression *large-scale behaviour is similar to that of negatively curved spaces*, and to develop analogues of many fruitful constructions associated to hyperbolic spaces to this more general class of spaces (and groups).

1. Introduction

We consider henceforth geodesic metric spaces, i.e., spaces where every pair of points is connected by a geodesic (a curve every sub-arc of which has length the distance between its endpoints). Our goal is to study *large-scale properties* of such metric spaces in particular those that have negative curvature in the large scale. Before making precise definitions, let us consider what this should mean.

There are (at least) two useful ways to make sense of this, but we shall mainly take the first point of view. This is to introduce an equivalence relation between metric spaces X and Y which corresponds to them being the same in the large scale.

The right relation to introduce is that of quasi-isometry. We recall below the definition.

Definition 1.1. Suppose X and Y are metric spaces. A function $f : X \to Y$ is said to be a (λ, ϵ)-quasi-isometric embedding if

$$\forall p, q \in X,\ \frac{1}{\lambda} d_X(p,q) - \epsilon < d_Y(f(p), f(q)) < \lambda d_X(p,q) + \epsilon$$

. A (λ, ϵ)-quasi-isometry is a (λ, ϵ)-quasi-isometric embedding such that the ϵ-neighbourhood of $f(X)$ is all of Y.

This is the right relation for several reasons:

- For the universal cover of a compact space X, any two metrics that are pullbacks of metrics on X are quasi-isometric.
- Such a metric in fact depends up to quasi-isometry only on $\pi_1(X)$ and is quasi-isometric to the word metric corresponding to any finite system of generators.

Date: January 10, 2005.

- In particular a word-metric on a finitely generated group is independent up to quasi-isometry on the choice of generators.

Thus, we can speak of the large scale behaviour of a group up to quasi-isometry as an intrinsic property of the group. For instance the statement that in the large-scale $\mathbb{Z}^2$ is the same as $\mathbb{R}^2$ now makes sense.

Thus, we are concerned with a category whose objects are equivalence classes of metric spaces. The appropriate morphisms are *coarse Lipschitz maps*, with isomorphisms being quasi-isometries. All the notions we define henceforth must be intrinsic to the quasi-isometry class of a metric space.

The second fruitful approach, which we do not take here, is to re-scale the metric multiplying all the distances by ϵ and consider the limit as ϵ goes to zero. This amounts to looking at the space from a very great distance, so that everything looks small and the eye cannot distinguish fine features of the landscape. The limit one considers is the so called Gromov-Hausdorff limit. We do not enter into details here, but confine ourselves to a warning that this limit is by no means uniquely defined - one rather requires that every limit of subsequences has an appropriate property.

Models of negatively-curved spaces are the hyperbolic plane $\mathbb{H}$, $CAT(k)$ spaces with $k < 0$ and trees T (which are $CAT(-\infty)$ spaces). Thus we seek a characterisation of these examples that is invariant under quasi-isometry. There are several (equivalent) such characterisations we can use.

Remark 1.1. As we re-scale the metric on a negatively curved space (say a $CAT(k)$ space or a manifold with pinched negative sectional curvature), the curvature goes to $-\infty$. Hence the limit described above is always a so-called $\mathbb{R}$-tree, a generalisation of the notion of a tree. Thus in many ways trees can be regarded as our sole model spaces.

2. The definition of hyperbolicity

Of the many possible characterisations, we take as our definition of δ-hyperbolicity the property that triangles are *slim*.

Definition 2.1. Let $\delta > 0$. A geodesic triangle in a metric space X is said to be δ-slim if each of its sides is contained in the δ-neighbourhood of the union of the other two sides. A geodesic space X is said to be δ-hyperbolic if every triangle in X is δ-slim.

Proposition 2.1. *If $k < 0$ then every $CAT(k)$ space is δ-hyperbolic, where δ depends only on k.*

Proof. This follows immediately from the definition of $CAT(k)$ spaces once we show that the hyperbolic plane $\mathbb{H}$ is δ-hyperbolic for some δ. Suppose that some triangle in $\mathbb{H}$ is not δ-slim, then it contains a semi-circle of radius δ. As any triangle has area at most π, we can deduce an upper bound on δ. □

We leave the proof of the following proposition as an exercise to the reader.

Proposition 2.2. *Any tree T is 0-hyperbolic.*

3. Quasi-geodesics

Geodesics play a central role in Riemannian geometry (as do lines in Euclidean geometry). However, we are studying spaces only up to quasi-isometry, and a quasi-isometry does not preserve geodesics. In particular, it is not clear form the above definition that δ-hyperbolicity depends only on the quasi-isometry class.

Thus, we need to replace geodesics by the more robust notion of a quasi-geodesic.

Definition 3.1. A (λ, ϵ)-quasi-geodesic in X is a (λ, ϵ)-quasi-isometric embedding of an interval or $\mathbb{R}$ into X.

Observe that a geodesic is just an isometric embedding of $\mathbb{R}$ or an interval. It is clear from the definition that quasi-isometries preserve the set of quasi-geodesics.

That this is indeed a useful substitute for geodesics is due to deep properties of negative curvature. One should think of a quasi-geodesic as being a wiggly line near a geodesic. We shall see that this is indeed the correct picture in δ-hyperbolic spaces.

Theorem 3.1. *For all $\delta > 0$, $\lambda \geq 1$ and $\epsilon \geq 0$ there exists a constant $R = R(\delta, \lambda, \epsilon)$ such that the Hausdorff distance between any (λ, ϵ)-quasi-geodesic $c : [p, q] \to X$ in a δ-hyperbolic space X and any geodesic α joining its endpoints is at most R.*

Before entering into the proof, we consider a couple of examples in Euclidean space to show that this is by no means true in the absence of negative curvature.

First, observe that a semi-circle is a (λ, ϵ)- quasi-geodesic with the constants being the same at all scales. Similarly, a sine-curve is a quasi-geodesic, with the constants depending only on the *ratio* of the amplitude to the wave-length. These are examples where there is no R depending only on the constants λ and ϵ.

Next, consider an infinite sine-curve whose amplitude and wavelength grow in the same proportion. This is a quasi-geodesic, but is not close to any geodesic. Another such example can be constructed using an infinite spiral. We shall see in contrast, as an easy consequence of the above, that in δ-hyperbolic spaces even infinite quasi-geodesics are close to geodesics.

Furthermore, in the case of the standard Cayley graph for $\mathbb{Z}^2$, not even geodesics satisfy the above theorem.

There are some immediate consequences of the theorem.

Corollary 3.2. *A geodesic metric space is δ-hyperbolic for some δ iff every quasi-geodesic triangle is slim.*

Corollary 3.3. *Any space quasi-isometrically embedded in a hyperbolic space is hyperbolic.*

Here and henceforth by a hyperbolic space we mean δ-hyperbolic for some δ, where different spaces may have differing values of δ.

We shall now prove the above theorem. The proof is intricate and will be accomplished in three steps. Firstly, we show that, given a quasi-geodesic and a geodesic as above, the quasi-geodesic can be replaced by one that is continuous. Next, we show that every point in the geodesic is close to some point in the quasi-geodesic. Finally, we prove the result.

Lemma 3.4. *Let X be a geodesic metric space. Given any (λ, ϵ)-quasi-geodesic $c : [p, q] \to X$, one can find a continuous $(\lambda, \epsilon' = \lambda + \epsilon)$-quasi-geodesic $c' : [a, b] \to X$ with the same endpoints such that:*

(1) *The Hausdorff distance between the images of* c *and* c' *is at most* $\lambda + \epsilon$
(2) *For all* $t, t' \in [a, b]$, $l(c'|_{[t,t']}) \leq k_1 d(c(t), c(t')) + k_2$, *where* k_1 *and* k_2 *depend only on* λ *and* ϵ

Proof. One takes c' to consist of geodesic segments joining the points of $c(\Sigma)$, where $\Sigma = \{a, b\} \cup (\mathbb{Z} \cup [p, q])$. The result follows by a straightforward calculation. □

Next we have the following lemma.

Lemma 3.5. *Let* X *be a* δ*-hyperbolic space. Let* c *be a continuous rectifiable path in* X*. If* γ *is a geodesic segment connecting the end-points of* c*, then for every* $x \in \gamma$*,*

$$d(x, im(c)) \leq |log_2(l(c))| + 1$$

Proof. Let $x \in \gamma$. We can assume that $c : [0, 1] \to X$ and is parametrised by arc-length.

Consider a geodesic triangle with vertices $c(0)$, $c(1/2)$ and $c(1)$. By the slim triangle condition, x is within a distance δ from some point y_1 from one of the other two sides of this triangle. Replacing γ by this side and x by y_1, we can proceed by induction to get a sequence of points y_n at distance at most $n\delta$ from x. At each stage $l(c)$ is halved, so that after at most $|log_2(l(c))|$ stages it becomes less than 1. We can now find y within a distance 1 of y_n. □

We can apply this to our quasi-geodesic. However it is still less than we need, as the bound depends on the length, not just the constants of the quasi-geodesic. Thus, we need a uniform upper bound on D such that some point in $x \in \alpha$ has distance at least D from c. Here and henceforth, we assume that we have replaced c by c'.

Pick x and D as above that are maximal. Let y and z be points of α at distance $2D$ from x on the the two sides of x (if there are no such points, pick the end-point instead. Then there are points y' and z' on c that are a distance at most D from y and z.

Consider the path from y to z consisting of geodesic segments from y to y' and z to z' and a subsegment of c. By hypothesis and the triangle inequality this has length at most $6Dk_1 + k_2 + 2D$. Thus, by the lemma, as the point x has distance D from this segment,

$$D - 1 \leq \delta log_2(6Dk_1 + k_2 + 2D)$$

from which a universal upper bound D_0 follows.

Finally, we show that if $R = D_0(1 + k_1) + k_2/2$, we show that every point in c is contained in a R neighbourhood of α. Suppose not, let $[a', b']$ be a maximal interval whose image lies outside this neighbourhood. By connectedness, we can find $w \in \alpha$, $t \in [a, a']$ and $t' \in [b', b]$ such that w has distance at most D_0 from $c(t)$ and $c(t')$, so that the distance between these is at most $2D_0$. One immediately deduces a contradiction to the hypothesis that c is a quasi-geodesic. □

3.1. k**-local geodesics.** We see next that there is a local criterion for recognising quasi-geodesics in a δ-hyperbolic space.

Definition 3.2. Let X be a geodesic metric space. A path $c : [a, b] \to X$ is said to be a k-local geodesic if $d(c(t), c(t')) = |t - t'|$ if $|t - t'| < k$

Theorem 3.6. *Let X be a δ-hyperbolic space and c a k-local geodesic with $k > 8\delta$. Then:*

1. *$im(c)$ is contained in the 2δ-neighbourhood of any geodesic $[c(a), c(b)]$ joining its endpoints.*
2. *$[c(a), c(b)]$ is contained in the 3δ-neighbourhood of $im(c)$.*
3. *c is a quasi-geodesic (with constants depending on k and δ).*

The proof of this is similar to those of the last section. The reader is referred to Bridson-Haefliger for details.

4. Reformulations of hyperbolicity

We shall now consider two equivalent formulations of hyperbolicity, the second of which was Gromov's original definition. These can best be thought of as saying that δ-hyperbolic spaces are close to trees, or rather that finite polygons in such spaces can be well-approximated by metric trees.

4.1. **Thin triangles.** Our first condition says that triangles can be approximated by *tripods*, metric trees with at most three edges and at most one vertex of degree greater than one. Given three non-negative numbers a, b and c there is a unique such tree with these as edge-lengths.

Now suppose $d(x, y)$, $d(y, z)$ and $d(z, x)$ are the lengths of three sides of a triangle in a geodesic metric space X (in fact any three numbers satisfying the triangle inequalities), there exist unique non-negative numbers a, b and c such that $d(x, y) = a + b$, $d(y, z) = b + c$ and $d(z, x) = c + a$.

Given a geodesic triangle $\Delta = \Delta(x, y, z)$. Let $T_\Delta = T(a, b, c)$. There is a natural isometry χ_Δ from the sides of the triangle to the tripod. Denote the vertices of the tripod by v_x, v_y and v_z and let o_Δ denote the central vertex of the tripod. The points of $\chi^{-1}(o_\Delta)$ are called the *internal points*.

The diameters of the fibres of χ are a natural measure of the thinness of the triangle. There are two measures of this.

Definition 4.1. A geodesic triangle Δ is said to be δ-thin if $\chi(p) = \chi(q) \Rightarrow d(p, q) \leq \delta$. The diameter of $\chi^{-1}(o_\Delta)$ is called the insize of Δ.

Proposition 4.1. *For a geodesic metric space X, the following conditions are equivalent.*

1. *X is δ_0-hyperbolic for some δ_0*
2. *There exists δ_1 such that every triangle in X is δ_1-thin.*
3. *There exists δ_2 such that every triangle in X has insize at most δ_2.*

Proof. It is clear that 2 implies 1. We next show that 1 implies 3. Let i_x denote the internal point on the side opposite x and similarly for i_y and i_z. By hypothesis, there is a point p in (say) $[x, y]$ that is a distance at most δ_0 from i_x. By the triangle inequality, $|d(y, p) - d(y, i_x)| \leq \delta_0$. As $d(y, i_x) = d(y, i_z)$ it follows that $d(p, i_z) \leq \delta_0$. Thus, $d(i_x, i_z) \leq 2\delta_0$. A similar argument with the other pairs completes the proof.

Finally, assume 3 holds. Consider a geodesic triangle and let $\{p, q\}$ be a fibre of χ. Without loss of generality we assume that $\chi(p)$ lies on the branch of the tripod ending in v_y and that $p \in [x, y]$. We show that 1 holds by constructing a geodesic triangle with p as an internal point. As $d(y, p) = d(y, q)$, it will follow that q is also an internal point, giving the result.

To construct the triangle, we look at triangles with vertices z, y and w, where w moves continuously on $[x, y]$. By our assumptions, when $w = x$, the internal vertex on $[x, y]$ is i_z which lies between p and x.

On the other hand, if w is close to y, the internal point lies between p and y. Using the intermediate value theorem, we get the desired triangle. □

4.2. **The Gromov Product.** We now turn to Gromov's original definition of hyperbolicity. Like the previous condition, it is best to think of this in terms of trees.

Definition 4.2. Let X be a geodesic metric space and let $x \in X$. The Gromov product of $y, z \in X$ with respect to x is defined to be

$$(y \cdot z)_x = \frac{1}{2}(d(x, y) + d(y, z) - d(z, x))$$

If X is a tree, this is the distance between x and the segment $[y, z]$. More generally, it is the corresponding length in the comparison tripod. As a consequence, if Δ is δ-thin then $|d(x, [y, z]) - (y \cdot z)_x| \leq \delta$.

Definition 4.3. A metric space is X is said to be (δ)-hyperbolic if

$$(x \cdot y)_w \geq min((y \cdot z)_w, (x \cdot z)_w) - \delta$$

for all $x, y, z, w \in X$.

By considering cases, it is an easy (*but important*) exercise to check that if X is a tree then the above holds with $\delta = 0$. Note that the above definition does not require X to be a geodesic metric space.

The condition can be formulated more symmetrically as the following *four-point condition.*

$$d(x, w) + d(y, z) \leq max(d(x, y) + d(z, w), d(x, z) + d(y, w)) + 2\delta$$

Once more it is useful to think of this in terms of trees and see that it holds with $\delta = 0$. The extremal case is when we have a six vertex tree with two trivalent vertices with an edge joining them, so that the pairs of points on the left hand side of the above inequality are on the opposite sides of this edge. The middle edge gets counted twice and every other edge once on both sides of the above.

Proposition 4.2. *Let X be a geodesic metric space. Then X is hyperbolic iff it is (δ)-hyperbolic for some δ.*

Proof. Suppose X is hyperbolic, that is, every triangle is δ-thin. Then we deduce (δ')-hyperbolicity by comparing the points x, y, z and w to points on a tree. This consists of two triads glued together along a common edge.

By considering cases we see that the tree is (0)-hyperbolic, and passing to X changes each side of the above inequality by at most 2δ (as our tree is constructed from two triads).

Conversely, we shall show that the insize of a triangle with vertices x, y and z is bounded. Consider the four-point for these points together with i_x. This implies that $d(x, i_x) + d(y, z) < P/2 + \delta$, where P is the perimeter of the triangle (one has equality in the case of a tree). Since $d(x, i_x) + d(y, z) = P/2$ it follows that $|d(x, i_x) - d(x, i_z)| < 2\delta$. Similarly, $|d(z, i_x) - d(z, i_z)| < 2\delta$.

Now using the four-point condition for the points x, z, i_x and i_z gives an upper bound on the distance between i_x and i_z.

□

5. The Linear isoperimetric inequality

An important property of hyperbolic spaces is that they satisfy linear isoperimetric inequalities, that is, any curve in a hyperbolic space *coarsely* bounds a disc whose area is a linear function of the length of the curve. We begin by defining coarse notion of the area of a bounding disc.

Definition 5.1. Let X be a metric space and $c : S^1 \to X$ be a rectifiable loop. A coarse ϵ-filling of c is a triangulation of D^2 together with a map $\Phi : D^2 \to X$ which restricts to c such that the diameter of the image of each triangle is at most ϵ. The ϵ-area of the filling is defined to be the number of triangles in the triangulation of the disc.

Theorem 5.1. *Let X be a δ-hyperbolic space. Then there are positive constants ϵ, A and B such that every rectifiable loop $c : S^1 \to X$ admits a coarse ϵ-filling with ϵ-area at most $Al(c) + B$.*

The first step in the proof is to replace X by a quasi-isometric space which is a metric tree with sides of integral length. To do this find a maximal set S in X so that the distance between every pair of points is at least 1. Join pairs of points in S whose distance is at most 3 by an edge whose length is the nearest integer to their distance. The map that takes each point in X to a point of minimal distance from it in S gives the desired quasi-isometry.

As quasi-isometries distort distances by a bounded linear amount, one can replace X in the above by the metric tree just constructed. For such a tree, we construct an ϵ-filling satisfying a linear isoperimetric inequality with $\epsilon = 16\delta$.

The heart of the proof is in the following lemma.

Lemma 5.2. *Given any locally-injective loop $c : [0,1] \to X$ beginning at a vertex, one can find $s, t \in [0,1]$ such that $c(s)$ and $c(t)$ are vertices of X, $d(c(s), c(t)) \leq l(c|_{[s,t]}) - 1$ and $d(c(s), c(t)) + l(c|_{[s,t]}) - 1 \leq 16\delta$*

Proof. As X is hyperbolic, c cannot be a k-local geodesic for $k = 8\delta + 1/2$, and hence we can a sub-arc whose length is less than $k = 8\delta + 1/2$ which is not a geodesic. Let $c(s)$ and $c(t)$ be the first and last vertices of this sub-arc. Then these satisfy the conclusion of the lemma, as it is not a geodesic and both the distance betzeen its endopints and the length of the segment are integers, so must differ by at least 1. □

Notice that such a statement does not hold in Euclidean space, as can be seen by considering very large circles.

Now the isoperimetric inequality can be proved by induction. By choosing constants appropriately, we can ensure this is satisfied for curves of length at most 16δ.

Now suppose a loop c is given. If c is not locally injective, we can find cancelling adjacent segments, and apply induction to the curve with these deleted. Else we use the lemma to find s and t, and consider the loops with $c|_{[s,t]}$ replaced by a geodesic, and the loop consisting of $c|_{[s,t]}$ and the geodesic. We can apply the induction hypothesis to each of these to get an ϵ-filling, and glue these together to get the required filling for c. □

6. The Gromov boundary ∂X

The hyperbolic plane viewed as the interior of the unit disc in $\mathbb{C}$ has a natural boundary, namely the unit circle. This is not an accident of the representation, but is intrinsically associated to the hyperbolic plane.

To construct the boundary, one considers the set of rays and defines two rays to be asymptotic if the Hausdorff distance between their boundaries is finite (i.e., they are parallel). We regard asymptotic rays as equivalent, and the set of equivalence classes of rays is the boundary ∂X.

This construction associates a set ∂X to any δ-hyperbolic space, in fact to any geodesic metric space. However, we would like this to be a construction invariant under quasi-isometries. This allows us, for instance, to associate a boundary to a group.

To show quasi-isometry invariance for δ-hyperbolic spaces, we make an analogous construction using quasi-geodesics in place of geodesics. Thus, let $\partial_q X$ be the set of equivalence classes of quasi-geodesic rays, where again quasi-geodesic rays are equivalent if they are asymptotic, i.e., their images are a bounded Hausdorff distance apart.

As quasi-geodesics are close to geodesics, it readily follows that $\partial_q X = \partial X$. As $\partial_q X$ is obviously invariant under quasi-isometries, so is ∂X

Remark 6.1. As we have seen, in general quasi-geodesics are not necessarily close to geodesics, so we expect ∂X to be smaller than $\partial_q X$. In fact ∂X is not in general a quasi-isometry invariant.

We now turn to an alternative construction of the boundary using Gromov products. Fix a base-point $p \in X$. We consider sequences of points in X and define when such a sequence converges to a point on the boundary. We then impose an equivalence relation on such sequences corresponding to when two such sequences converge to the same point.

Definition 6.1. A sequence $x_n \in X$ is said to admissible if $(x_i \cdot x\boldsymbol{j})_p \to \infty$ when $i, j \to \infty$. Two admissible sequences $x_n \in X$ and $y_n \in X$ are said to be equivalent if $(x_i \cdot x\boldsymbol{j})_p \to \infty$ when $i, j \to \infty$

It is easy to see that the equivalence classes of admissible sequences under this relation correspond to points of ∂X. Furthermore, one can see that the Gromov product extends to one on $X \cup \partial X$, which is finite except for the product of a point in ∂X and itself. This is defined by

$$(x \cdot y)_p = sup\ liminf_{i,j \to \infty} (x_i \cdot y_j)_p$$

where the supremum is taken over all sequences with limits x and y.

7. The topology of ∂X

Using the Gromov product defined in the previous section, it is easy to define a topology on ∂X such that the space $X \cup \partial X$ is compact. Basic neighbourhoods of points in X in this topology are the same as those for the topology on X. For a point $z \in \partial X$, we take basic neighbourhoods to be $\{x \in X \cup \partial X : (x \cdot z)_p > R\}$.

The topology on the boundary is metrisable. One gets a pseudo-metric by taking $d(x, y) = e^{-(x \cdot y)_p}$ for points x, y in the boundary. The topology does not depend on p, though the metric does.

In a 0-hyperbolic space, this is in fact an ultra-metric, but in general it does not satisfy the triangle inequality, though we can control the extent to which it fails using the δ-hyperbolicity condition on Gromov products. To make this into a genuine metric, replace the distance between x and y by the infimum of the sum of distances betzeen succesive pairs of points for a sequence of points from x to y.

The metrics we have defined are the *visual metrics* on the boundary.

Department of mathematics, SUNY at Stony Brook, Stony Brook, NY 11794
E-mail address: gadgil@math.sunysb.edu