Contents

1 Foundations of geometric group theory

1.1 Word metrics and Cayley graphs

Let S be a finite generating set for a group G. We assume S is *symmetric*, i.e., if $s \in S$, then $s^{-1} \in S$.

Definition 1.1. A *word* in S is a concatenation of elements of S: $W = s_1 s_2 \dots s_n$ where $s_i \in S$. We let $\|W\|$ denote the length of the word, so here $\|W\| = n$. If two words W, U are letter-for-letter equivalent, we write $W \equiv U$. If W, U represent the same element of the group G, we write $W =_G U$. If $g \in G$, write $|g|_S$ for the length of the shortest word in S that represents g in G. Given $g, h \in G$, let $d_S(g,h) = |g^{-1}h|_S$. The metric d_S is the *word metric* on G with respect to S.

Definition 1.2. The *Cayley graph* of G with respect to S, which we denote by $\Gamma(G,S)$ is the graph whose vertex set is $\{g \in G\}$ with a (directed) edge from a vertex g to a vertex h whenever $h = gs$ for some $s \in S$.

Each edge has an orientation and a label. But if $gs = h$, then $g = hs^{-1}$ as well, so there is also a directed edge from h to g labeled by s^{-1}. This is rather cumbersome, so we typically don't draw the edge labeled by s^{-1}. We just keep in mind that if we traverse an edge labeled by s going in the opposite direction as its orientation, we think of that edge as labeled by s^{-1}.

Given a path p in $\Gamma(G,S)$, let $\mathbf{Lab}(p)$ be its label in $\Gamma(G,S)$. We denote the initial point of p by p_- and the terminal point of p by p_+.

Remark 1.3. d_S is a metric on the vertex set of $\Gamma(G,S)$. Identifying each edge with the unit interval $[0,1]$ extends this metric to the whole graph.

Exercises 1.4. *Draw the following Cayley graphs.*

1. $\Gamma(\mathbb{Z}/4\mathbb{Z}, \{1\})$
2. $\Gamma(\mathbb{Z}, \{1\})$
3. $\Gamma(\mathbb{Z}, \{2,3\})$
4. $\Gamma(\mathbb{Z}^2, \{a = (1,0), b = (0,1)\})$
5. $\Gamma(\mathbb{F}_2, \{a,b\})$, *where* $\mathbb{F}_2 = \langle a, b\rangle$ *(i.e., the group with two generators and no relations)*
6. $\Gamma(BS(1,2), \{a,t\})$, *where* $BS(1,2) = \langle a, t \mid tat^{-1} = a^2\rangle$.

1.2 Metric spaces

Let (X, d) denote a metric space X with its metric d. If the metric is clear, we may just write X.

If $x \in X$ and $r \geq 0$, then $B_R(x) = \{y \in X \mid d(x, y) \leq r\}$ is the (closed) ball of radius r about x. If $A \subseteq X$, then we write $\mathcal{N}_r(A) = \{y \in X \mid d(y, a) \leq r\} = A^{+r}$ for the (closed) r–neighborhood of A in X.

Definition 1.5. A *path* is X is a continuous map $p\colon [a, b] \to X$ for some $[a, b] \subseteq \mathbb{R}$. By an abuse of notation, we also call the image of this map p. Then $p_- = p(a)$ and $p_+ = p(b)$. A *ray* is a continuous map $p\colon [a, \infty) \to X$, and a *bi-infinite path* is a continuous map $p\colon (-\infty, \infty) \to X$.

Given a path p, the *length* of p is

$$\ell(p) := \sup_{a \leq t_1 \leq t_2 \leq \cdots \leq t_n \leq b} \sum_{i=1}^{n-1} d(p(t_i), p(t_{i+1})), \tag{1}$$

where the supremum is taken over all n and all choices of $t_1, \ldots, t_n$. This supremum is often infinite. We typically consider *rectifiable* paths, which are those with $\ell(p) < \infty$.

Example 1.6 (From BH I.1 Ex. 1.19). Let $X = [0, 1]$, and let $0 = t_0 < t_1 < \cdots < t_n < \cdots$ be an infinite sequence in $[0, 1]$ converging to 1. let $c\colon [0, 1] \to X$ be any path such that $c(0) = 0$, $c(t_n) = \sum_{k=1}^{n} \frac{(-1)^{k+1}}{k}$. Then c is not rectifiable because its length is bounded below by the harmonic series $\sum_{i=1}^{\infty} \frac{1}{k}$, which diverges to infinity.

Definition 1.7. A path is *geodesic* if $\ell(p) = d(p_-, p_+)$. A ray is *geodesic* if for all $a \leq s \leq t < \infty$, $\ell(p|_{[s,t]}) = d(p(s), p(t))$, and similarly for a bi-infinite path. A metric space is *geodesic* if every pair of points can be connected by a geodesic.

A weaker condition one can put on a metric space is that it is a *length space*, which means that it is a path-connected space such that for all $x, y \in X$, $d(x, y) = \inf\{\ell(p) \mid p_- = x, p_+ = y\}$. For example, consider a space with two points, x and y, with countably many edges between them indexed by the positive integers so that the n-th edge has length $1 + \frac{1}{n}$. Here, $d(x, y) = \inf(1 + frac1n)$, but there is no path of length 1 between them. Most statements in this class will hold for length spaces, but the proofs tend to be a bit more technical, so we will mostly work with geodesic spaces.

Definition 1.8. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f\colon X \to Y$ is:

- an *isometry* if for all $x, y \in X$,
$$d_Y(f(x), f(y)) = d_X(x, y)$$
and f is surjective;
- a *biLipschitz equivalence* if there exists $\lambda \geq 1$ such that for all $x, y \in X$,
$$\frac{1}{\lambda} d_X(x, y) \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y)$$
and f is surjective;
- a (λ, c)–*quasi-isometry* if there exists $\lambda \geq 1$ and $c \geq 0$ such that for all $x, y \in X$
$$\frac{1}{\lambda} d_X(x, y) - c \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y) + c$$
and $f(X)^{+c} = Y$, i.e., f is *coarsely surjective.*

We write $X \sim_{Lip} Y$ or $X \sim_{qi} Y$ if there exists a biLipschitz equivalence or (λ, c)–quasi-isometry from X to Y, respectively.

If the assumption of (coarse) surjectivity is dropped, then f is a(n) isometric, biLipschitz, or quasi-isometric *embedding*, respectively. Note that quasi-isometries are not required to be continuous, though it will follow from the definitions that isometries and biLipschitz equivalences are continuous.

Exercises 1.9.

1. *Show that $\sim_{Lip}$ and $\sim_{qi}$ are equivalence relations.*
2. *Let X, Y be bounded metric spaces. Show that $X \sim_{qi} Y$.*
3. *Show that $\Gamma(\mathbb{Z}, \{1\})$ is quasi-isometric to the real line $\mathbb{R}$.*
4. *Show that $\Gamma(\mathbb{Z}, \{1\}) \sim_{qi} \Gamma(\mathbb{Z}, \{2, 3\})$.*
5. *In general, show that if $S, T \subseteq G$ are two finite generating sets, then $(G, d_S) \sim_{Lip} (G, d_T)$, and so $\Gamma(G, S) \sim_{qi} \Gamma(G, T)$.*

1.3 Group actions

Let G act on a metric space X, that is, there exists a map $\varphi\colon G \times X \to X$ such that $\varphi(gh, x) = \varphi(g, \varphi(h, x))$. We write $G \curvearrowright X$, and, to save space, we often drop the φ notation simply write gx (or $g.x$ or $g \cdot x$) for $\varphi(g, x)$. In this class, we will ALWAYS assume the action is by isometries, that is, for all $g \in G$, $d_X(x, y) = d_X(gx, gy)$.

There is a bijection between actions of G on X and homomorphims $\rho\colon G \to \mathrm{Isom}(X)$, where $\mathrm{Isom}(X)$ is the group of isometries of the space X.

Definition 1.10. The action is *faithful* if ρ is injective, i.e., for all $g \in G$, there exists $x \in X$ such that $gx \neq x$. The action is *free* if for all $x \in X$, $\mathrm{Stab}_G(x) = \{1\}$, where $\mathrm{Stab}_G(x) = \{g \in G \mid gx = x\}$, i.e., for all $x \in X$ for all $g \in G$, $gx \neq x$.

The action is *proper* if for all bounded subsets $B \subseteq X$,

$$\#\{g \in G \mid gB \cap B \neq \emptyset\} < \infty.$$

This is also called *metric properness.* There is a topological version where bounded is replaced by compact.

The action is *cobounded* (resp. *cocompact*) if there exists a cobounded (resp. cocompact) subset $K \subseteq X$ such that

$$\bigcup_{g \in G} gK = X.$$

Group actions can be used to prove algebraic results. Here is just one example:

Theorem 1.11. *Suppose G acts freely on $\mathbb{R}^n$ (where $\mathbb{R}^n$ has the Euclidean metric). Then G is torsion free.*

Proof. Exercise. □

[Hint: Use the fact that any finite set of points $\{x_1, \ldots, x_n\}$ in $\mathbb{R}^n$ has a unique *centroid*, that is, a point $w \in \mathbb{R}^n$ that minimizes the distance $\sum d_{\mathbb{R}^n}(w, x_i)$.]

The following is one of the the foundational results of geometric group theory. A *proper* metric space is one in which closed balls are compact.

Theorem 1.12 (Milnor–Swarz Lemma)**.** *Let G be a group acting properly and coboundedly on a propoer geodesic metric space X. Then G has a finite generating set S and $\Gamma(G, S) \sim_{qi} X$.*

Proof. Fix $x \in X$. Choose a bounded subset K of X so that $\bigcup_{g\in G} gK = X$. Since K is bounded, there exists a closed ball B of radius D such that $K \subseteq B$. Thus we have $\bigcup_{g\in G} gB = X$. Since X is proper, B is compact. Let

$$S = \{g \in G \setminus \{1\} \mid B \cap gB \neq \emptyset\}.$$

Since $G \curvearrowright X$ is proper, the set S is finite, and $S = S^{-1}$ by the definition of S.

Consider $\inf\{d(B, gB) \mid g \in G \setminus (S \cup \{1\})\}$. For some $g \in G \setminus (S \cup \{1\})$, the distance $d(B, gB)$ is a positive constant R, by the definition of S. The set H of elements of G such that $d(B, gB) \leq R$ is contained in

$$\{g \in G \mid gB(x, D+R) \cap B(x, D+R)\}$$

and hence is finite. Since

$$\inf\{d(B, gB) \mid g \in G \setminus (S \cup \{1\})\} = \inf\{d(B, gB) \mid g \in H \setminus (S \cup \{1\})\},$$

and the latter is taken over a finite set, we have that the infimum is realized by some $h_0 \in H \setminus (S \cup \{1\})$. For this h_0, we have that $d(B, hB) = 2L$ for some positive constant L. By definition, if $d(B, gB) < 2L$, then $g \in S$.

We first show that S generates G. Let $g \in G$ and consider a geodesic $[x, gx]$ in X. let $k = \left\lfloor \frac{d(x,gx)}{L} \right\rfloor$. Then there exists a finite sequence a points along the geodesic $x = x_0, x_1, \dots x_{k+1} = gx$ such that $d(x_i, x_{i+1}) \leq L$ for every $i \in \{0, \dots, k\}$. For each such i, let $h_i \in G$ be such that $x_i \in h_iB$. We take $h_0 = 1$ and $h_{k+1} = g$. As $d(B, h_i^{-1}h_{i+1}B) = d(h_iB, h_{i+1}B) \leq d(x_i, x_{i+1}) = L$, it follows that $h_i^{-1}h_{i+1} \in S$. Also,

$$h_0(h_0^{-1}h_1)(h_1^{-1}h_2)\cdots(h_k^{-1}h_{k+1}) = h_{k+1} = g,$$

and so $g \in \langle S \rangle$. Therefore the finite set S generates G.

It remains to show that $\Gamma(G, S)$ is quasi-isometric to X. let $f\colon G \to X$ be defined by $f(g) = gx$ (so f is the orbit map). Since $G \curvearrowright X$ is cobounded, this map is $2D$–coarsely surjective.

The argument above shows that

$$|g|_S \leq k + 1 \leq \frac{1}{L} d(x, gx) + 1.$$

For the other half of the inequality, let $|g|_S = m$, and let $w = s_1 s_2 \dots s_m$ be a word in S such that $w =_G g$. By the triangle inequality,

$$\begin{aligned}
d(x, gx) &= d(x, s_1 s_2 \dots s_m) \\
&\leq d(x, s_1 x) + d(s_1 x, s_1 s_2 x) + \cdots + d(s_1 \dots s_{m-1} x, s_1 \dots s_m x) \\
&= \sum_{i=1}^{m} d(x, s_i x) \\
&\leq 2Dm \\
&= 2D|g|_S.
\end{aligned}$$

Thus, we have for any $g \in G$:

$$L|g|_s - L \leq d(x, gx) \leq 2D|g|_S.$$

Note that since both the word metric and the metric on X are left-invariant (under the action of G), this suffices to prove the statement for any $h, g \in G$, rather than the pair $1, g \in G$ for which we have proven the result. □

Corollary 1.13. *The following follow immediately from the Milnor–Schwarz Lemma:*

1. *If G is finitely generated and H is a finite-index subgroup of G, then $H \sim_{qi} G$.*
2. *If $N \trianglelefteq G$ is a finite normal subgroup and G/N is finitely generated, then G is finitely generated and $G \sim_{qi} G/N$.*
3. *If M is a closed Riemannian manifold and $\tilde{M}$ is its universal cover, then $\pi_1(M) \sim_{qi} \tilde{M}$.*
4. *If G is a connected Lie group with left-invariant Riemannian metric and $\Gamma \leq G$ is a uniform lattice, then Γ is finitely generated and $\Gamma \sim_{qi} G$.*

2 Geometry of hyperbolic metric spaces

Let X be a metric space. Given two points $x, y \in X$, let $[x, y]$ be a geodesic from X to Y.

2.1 Four equivalent definitions of hyperbolicity

Assume X is geodesic. For any geodesic triangle T in X with vertices x, y, z the following hold.

1. There exist points $i_z \in [x, y], i_y \in [x, z], i_x \in [y, z]$ such that $a := d(x, i_z) = d(x, i_y)$, $b := d(y, i_z) = d(y, i_x)$ and $c := d(z, i_x) = d(z, i_y)$. See Figure 1. The points i_x, i_y, i_z are called *internal points* of T.

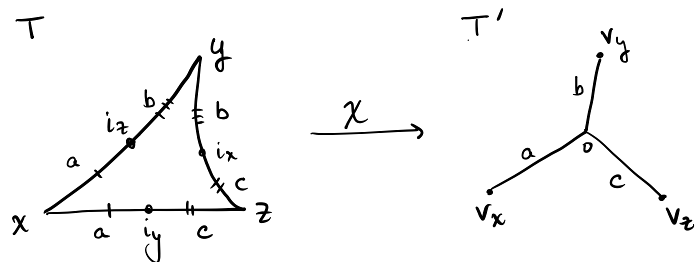

Figure 1: A geodesic triangle T in X and its comparison tripod T'.

2. Let T' be a *comparison tripod*, that is, a graph with 3 vertices of valence one labeled v_x, v_y, and v_z, and a central vertex of valence 3 labeled o, so that the lengths of the edges from v_x, v_y, and v_z to o are a, b, and c, respectively. The map $T \to T'$ which sends x, y, z to v_x, v_y, v_z respectively, extends uniquely to a map $\chi \colon T \to T'$ which restricts to an isometry on each side of T. This map sends i_x, i_y, i_z to o, and every other point in T' except v_x, v_y, and v_z will have a preimage in T consisting of exactly two points.

Definition 2.1 (Thin triangles)**.** Let $\delta \geq 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_1(\delta)$ if for any geodesic triangle T in X and any $a, b \in T$ which map to the same point in T', we have $d(a, b) \leq \delta$. The triangle T is called *δ–thin*.

Definition 2.2 (Insize)**.** Let $\delta \geq 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_2(\delta)$ if for any geodesic triangle T in X, $\mathrm{diam}(\{i_x, i_y, i_z\}) \leq \delta$. The quantity $\mathrm{diam}(\{i_x, i_y, i_z\})$ is called the *insize* of the triangle T.

The following is the most common definition of a (geodesic) hyperbolic metric space. It is attributed to Rips, and is sometimes called the Rips condition.

Definition 2.3 (Slim triangles)**.** Let $\delta \geq 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_3(\delta)$ if for any geodesic triangle T in X, we have $[x,y] \subseteq \mathcal{N}_\delta([x,z] \cup [y,z])$. Equivalently, given any point $a \in [x,y]$, there exists a point $b \in [x,z] \cup [y,z]$ such that $d(a,b) \leq \delta$.

Now let X be a (not necessarily geodesic) metric space and $x_0 \in X$.

Definition 2.4. The *Gromov product* of $x, y \in X$ with respect to x_0 is

$$(x \mid y)_{x_0} = \frac{1}{2}(d(x_0,x) + d(x_0,y) - d(x,y)).$$

If X is a tree, then $(x \mid y)_{x_0}$ measures the distance from x_0 to the geodesic $[x,y]$ (draw a picture of this!). Intuitively, $(x \mid y)_{x_0}$ will in general measure how long geodesics from x_0 to x and from x_0 to y "fellow travel." But note that we did not assume X was geodesic. Another way to think about the Gromov product is that it measures the failure of the triangle inequality to be an equality: in general we have that $d(x_0,x) + d(x_0,y) \geq d(x,y)$, and we call $d(x_0,x) + d(x_0,y) - d(x,y) \geq 0$ the *defect*. The quantity $(x \mid y)_{x_0}$ is half the defect.

The following is Gromov's original definition of a hyperbolic metric space.

Definition 2.5. Let $\delta \geq 0$. A metric spaces satisfies $\mathrm{Hyp}_4(\delta)$ if for any four points $x, y, z, w \in X$,

$$(x \mid y)_w \geq \min\{(x \mid z)_w, (y \mid z)_w\} - \delta.$$

Exercises 2.6. *If X is a tree, then X satisfies Hyp$_4(0)$. (We discussed this in class, but see if you can reconstruct the reasoning!)*

We can reformulate $\mathrm{Hyp}_4(\delta)$ as a more symmetric 4–*point condition*:

$$d(x,w) + d(y,z) \leq \max\{d(x,y) + d(z,w), d(x,z) + d(y,w)\} + 2\delta. \tag{2}$$

If you consider x, y, z, and w to be the vertices of a tetrahedron, then the three sums in (2) are the sums of the lengths of the pairs of opposite sides. In a tree, a tetrahedron is a subtree with 6 vertices: the vertices x, y, z, w have valence 1, and there are two vertices with valence three. In such a tetrahedron, you can check that there is a tie for the two largest sums of opposite sides, and so $d(x,w) + d(y,z) \leq \max\{d(x,y) + d(z,w), d(x,z) + d(y,w)$. The 4–point condition makes this coarse by requiring only that there is a *coarse* tie for the two largest sums.

Up to modifying δ by a constant multiple, all four conditions $\mathrm{Hyp}_i(\delta)$ are equivalent:

Theorem 2.7. *For any geodesic metric space X, any $\delta \geq 0$, and any geodesic triangle T, the following hold:*

(1) T is δ–thin $\implies$ T is δ–slim $\implies$ the insize of T is at most 2δ $\implies$ T is 2δ–thin; and

(2) the insize of T is at most δ $\implies$ Hyp$_4(\delta)$ $\implies$ the insize of T is at most 6δ.

Proof. We first prove (1). It is clear that the first implication holds.

For the second implication, let $x, y, z \in X$, and consider a geodesic triangle T. Suppose T is δ–slim. Then there exists a point p on (without loss of generality) $[x,y]$ at distance at most δ from i_x. By the triangle inequality, $|d(y,p) - d(y,i_x)| \leq \delta$. As $d(y,i_x) = d(y,i_z)$, we have $d(p,i_z) \leq \delta$. Thus $d(i_x,i_z) \leq 2\delta$. Similarly for the other pairs.

For the third implication, suppose p lies on $[y,z]$ is such that $d(y,p) < d(y,i_x)$. Then the fiber of $\chi\colon T \to T'$ containing p is $\{p,q\}$, where $q \in [y,x]$ and $d(y,p) = d(y,q)$. We will show that $d(p,q) \leq 2\delta$ by building a geodesic triangle with p, q as internal points.

let $c\colon [0,1] \to X$ be a monotone parametrization of $[y,z]$, and for each $t \in [0,1]$, consider a geodesic triangle $T_t = T(x,y,c(t))$, two of whose sides are $[y,x]$ and $c([0,t])$. The interal point of T_t on $c([0,t])$ varies continuously (though maybe not monotonically) as a function of t. At $t=0$, it is y, while at $t=1$ it is i_x, so for some $t \in [0,1]$ it is p. Since $d(y,p) = d(y,q)$, q must also be an internal point for T_t. As we assume the insize of all triangles is at most 2δ, we have $d(p,q) \le 2\delta$.

We now prove (2). For the first implication, assume that the insize of any geodesic triangle in X is at most δ. Given $x,y,z,w \in X$, we may assume without loss of generality that $S := d(x,z)+d(y,w) \le M := d(x,y)+d(z,w) \le L := d(x,w)+d(y,z)$. To show the 4–point condition holds, we must show that $L \le M + 2\delta$. Let $T = T(x,w,y)$ and $T' = T'(x,w,z)$ be geodesic triangles, and denote their internal points by i_x, i_w, i_y and i'_x, i'_w, i'_z, respectively. Consider the path from y to z which passes through i_x, i_y, i'_z, and i'_w, in that order. By the triangle inequality, we have

$$\begin{aligned} d(y,z) &\le d(y,i_x) + d(i_x,i_y) + d(i_y,i'_z) + d(i'_z,i'_w) + d(i'_w,z) \\ &\le d(y,i_x) + \delta + d(i_y,i'_z) + \delta + d(i'_w,z) \\ &= d(y,i_w) + d(i_y,i'_z) + d(z,i'_x) + 2\delta. \end{aligned}$$

We also have that

$$d(x,w) = d(x,i_y) + d(i'_z,w) - d(i'_z,i_y) = d(x,i_w) + d(i'_x,w) + d(i'_z,i_y).$$

Thus

$$\begin{aligned} L &\le d(y,i_w) + d(i_y,i'_z) + d(z,i'_x) + 2\delta + d(x,i_w) + d(i'_x,w) + d(i'_z,i_y) \\ &= (d(x,i_w) + d(i_w,y)) + (d(w,i'_x) + d(z,i'_x)) + 2\delta \\ &= M + 2\delta, \end{aligned}$$

as desired.

For the second implication, assume that X satisfies $\mathrm{Hyp}_4(\delta)$. Consider a triangle $T = T(x,y,z)$. We will apply the four-point condition to the points x,y,z, and $i_x \in [y,z]$. The largest sum of pairs of distances in the four-point condition will be $d(x,i_x)+d(y,z)$. To see this, notice that $2(d(x,i_x)+d(y,z)) = (d(x,i_x) + d(i_x,y)) + (d(x,i_x)+d(i_x,z)) + d(y,z)$ is greater than the perimeter P of T, while the other sums are $d(z,i_x)+d(x,y) = d(y,i_x)+d(x,z) = P/2$. Therefore the four-point condition implies that $d(x,i_x)+d(y,z) \le P/2 + 2\delta$. Since $d(y,z)+d(x,i_z)$ is also equal to $P/2$, we conclude that $|d(x,i_x) - d(x,i_z)| \le 2\delta$. Similarly, we have $|d(z,i_x) - d(z,i_z)| \le 2\delta$.

We now consider the four points $\{x,z,i_x,i_z\}$. In this case the three sums of pairs of distances are $d(x,z) = d(x,i_z)+d(z,i_x)$, $d(i_x,i_z)+d(x,z)$, and $d(x,i_x)+d(z,i_z)$. By the previous paragraph, the third sum is at most $d(x,i_z)+d(z,i_x)+4\delta = d(x,z)+4\delta$. Thus we see that the largest two pairs of sums are the last two listed above. Applying the four-point condition, we have $d(i_x,i_z) \le 6\delta$. Analogous arguments yield that $d(i_x,i_y) \le 6\delta$ and $d(i_y,i_z) \le 6\delta$, as well, completing the proof. □

Definition 2.8. We say that a metric space X is δ–hyperbolic (or Gromov hyperbolic) if it satisfies $\mathrm{Hyp}_4(\delta)$. If X is a geodesic metric space, we will also assume that δ has been chosen so that $\mathrm{Hyp}_1(\delta)$–$\mathrm{Hyp}_3(\delta)$ also hold.

Note: In a geodesic δ–hyperbolic metric space, the geometric interpretation of the Gromov product is coarsely the same as in the case of a tree: $(x \mid y)_{x_0}$ and $d(x_0,[x,y])$ will differ by at most δ.

Example 2.9. Examples of hyperbolic metric spaces:

1. If X is a bounded diameter metric space, then X is δ–hyperbolic with $\delta = \mathrm{diam}(X)$.

2. If X is a simplicial tree, then X is 0–hyperbolic.

3. The classical hyperbolic space $\mathbb{H}^2$ is δ–hyperbolic. (Exercise: Find δ! Hint: first show that it suffices to consider ideal triangles – those with all vertices at infinity. Then use the (isometric) action of $PSL_2(\mathbb{R})$ on $\mathbb{H}^2$ to show that it suffices to consider the ideal triangle with vertices at $-1, 1$, and ∞, in the upper half plane model.)

4. $\mathbb{H}^n$ for $n \geq 3$.

5. A space X is an *ultrametric space* if the following strong version of the triangle inequality holds for all $x, y, z \in X$: $d(x, y) \leq \max\{d(x, z), d(y, z)\}$. (In other words, a (geodesic) space is an ultrametric space is there is always a tie for the two longest sides of a triangle.) Ultrametric spaces often come up in relation to number theory; for example, the p–adic numbers $\mathbb{Q}_p$ with the p–adic metric space is an ultrametric space. All ultrametric spaces are 0–hyperbolic (this can be shown using the fourth definition of hyperbolicity).

6. An $\mathbb{R}$–tree is a generalization of a simplicial tree which can be defined as a space in which every geodesic triangle is a tripod. $\mathbb{R}$–trees are 0–hyperbolic.

7. Non-example: $\mathbb{R}^2$ (and more generally $\mathbb{R}^n$) is not hyperbolic for any δ. In fact, more is true: no δ–hyperbolic space can contain an isometrically copy of $\mathbb{R}^2$ (called a *flat*), or even a quasi-isometrically copy (called a *quasi-flat*).

Exercises 2.10. *Let X be a δ–hyperbolic metric space and $P = p_1 \dots p_n$ be a geodesic n–gon in X with $n \geq 3$. Let a be a point on p_i for some $1 \leq i \leq n$. Prove that there exists $j \neq i$ and $b \in p_j$ such that $d(a, b) \leq (n-2)\delta$. (In fact, $n-2$ can be replaced by* $\log_2(n)$*.)*

2.2 Quasigeodesic stability

We'd like to define a hyperbolic group as a finitely generated group so that every Cayley graph is δ–hyperbolic for some δ. However, this is a cumbersome (and impractical) definition, as we'd have to check it for *every single* finite generating set of the group. What would be nice is if the property of a metric space being δ–hyperbolic was a quasi-isometry invariant (in some sense which we'll make precise later). The most straightforward definitions of hyperbolicity involve geodesic triangles, and the problem is when you only look at things up to quasi-isometry, you lose control of how geodesics behave. In particular, the image of a geodesic under a quasi-isometry is not necessarily a geodesic any more. This motivates the notion of a *quasi-geodesic.*

Definition 2.11. A (λ, ϵ)–*quasi-geodesic* in a metric space X is a (λ, ϵ)–quasi-isometric embedding $c \colon I \to X$, where $I \subset \mathbb{R}$ is connected. More explicitly, we have

$$\frac{1}{\lambda}|t - t'| - \epsilon \leq d(c(t), c(t')) \leq \lambda|t - t'| + \epsilon$$

for all $t, t' \in I$. If $I = [a, b]$, then $c(a)$ and $c(b)$ are called the *endpoints* of c. If $I = [0, \infty)$, then c is called a *quasi-geodesic ray.*

We typically abuse notation and call the image of c the quasi-geodesic c. Note that since quasi-isometries are not necessarily continuous, neither are quasi-geodesics. In particular, a quasi-geodesic could just be a discrete set of points. However, it is shown in Bridson–Haefliger [**?**, Lemma III.H.1.11] that in any geodesic metric space, all quasi-geodesics are within Hausdorff distance at most $\lambda + \epsilon$ of a *tame* quasi-geodesic, that, is, a continuous quasi-geodesic. Because of this, when we work in geodesic metric spaces, we often consider only continuous quasi-geodesics.

(The *Hausdorff distance* between closed subsets A and B of a metric space X is the infimum of all ϵ such that $A \subseteq B^{+\epsilon}$ and $B \subseteq A^{+\epsilon}$. We write $d_{Haus}(A, B)$.)

Lemma 2.12. *If $c\colon I \to X$ is a continuous (λ, ϵ)–quasi-geodesic in a metric space X, then for any $t, t' \in I$, we have*

$$\ell(c|_{[t,t']}) \leq \lambda d(c(t), c(t')) + \epsilon,$$

where $\ell(c|_{[t,t']})$ denotes the length of the subpath of c from $c(t)$ to $c(t')$.

In general, quasi-geodesics in metric spaces can look nothing like geodesics. For example, the map $[0, \infty) \to \mathbb{R}^2$ given in polar coordinates by $t \mapsto (t, \log(1 + t))$ is a quasi-geodesic ray.

Lemma 2.13. *Let X be a δ–hyperbolic geodesic metric space. Let c be a continuous (rectifiable) path in X. If $[p, q]$ is a geodesic connecting the endpoints of c, then for every point $x \in [p, q]$, there exists a point $y \in c$ such that*

$$d(x, b) \leq \delta|\log_2(\ell(c))| + 1.$$

Proof. Exercise! ☐

There are two very important properties of quasi-geodesics in hyperbolic metric spaces: a local-to-global property, and the Morse lemma. We begin with the Morse lemma.

Theorem 2.14 (Morse Lemma; Quasi-geoesic stability)**.** *For all $\delta \geq 0$, $\lambda \geq 1$, and $\varepsilon \geq 0$, there exists a constant $R(\delta, \lambda, \varepsilon) \geq 0$ such that the following holds. If X is a δ–hyperbolic metric space, c is a (λ, ε)–quasi-geodesic in X, and $[p, q]$ is a geodesic connecting the endpoints of c, then*

$$d_{Haus}(c, [p, q]) \leq R.$$

Before giving the proof, we note several consequences of the Morse lemma.

Corollary 2.15. *For every $\lambda \geq 1$, $\varepsilon \geq 0$, and $\delta \geq 0$, there is a constant $M = M(\delta, \lambda, \varepsilon)$ such that a geodesic metric space X is δ–hyperbolic if and only if every (λ, ε)–quasigeodesic triangle in X is M–slim.*

Corollary 2.16. *Let X, X' be geodesic metric spaces and $f\colon X \to X'$ be a (λ, ε)–quasi-isometric embedding. If X' is δ–hyperbolic, then X is δ'–hyperbolic for some $\delta' = \delta'(\delta, \lambda, \varepsilon)$.*

Exercises 2.17. *Give careful proofs of both corollaries.*

A special case of the above corollary says that if $X \sim_{qi} X'$ and X is δ–hyperbolic, then X' is δ'–hyperbolic. This is what we mean by hyperbolicity being a quasi-isometry invariant.

We now prove the Morse Lemma.

Proof of the Morse Lemma. We may assume c is tame. Let $D = \sup\{d(x, c) \mid x \in [p, q]\}$, and let $x_0 \in [p, q]$ be the point where the supremum is attained (the point x_0 exists because the interval $[p, q]$ is compact). Then the open ball of radius D about x_0 is disjoint from c.

Let $y \in [p, x_0]$ be such that $d(y, x_0) = 2D$, and similarly for $z \in [x_0, q]$ (if such points do not exist, we take $y = p$ or $z = q$). Fix $y', z' \in c$ with $d(y, y') \leq D$ and $d(z, z') \leq D$. Choose geodesics $[y, y']$ and $[z, z']$, and let γ be the path from y to z formed by concatenating $[y, y']$, the subpath of c from y' to z', and $[z', z]$. Then by construction, γ lies outside of $B(x_0, D)$.

We have

$$\ell(\gamma) \leq 2D + \ell(c|_{[y',z']}) \leq 2D + \lambda d(y', z') + \varepsilon \leq 2D + \lambda \cdot 6D + \varepsilon.$$

Since $d(x_0, \gamma) = D$, we have $D \leq \delta|\log_2(\ell(\gamma))| + 1$, and so

$$D - 1 \leq \delta \log(2D + 6D\lambda + \varepsilon) + 1. \tag{3}$$

Since the left-hand side is linear in D while the right-hand side is logarithmic in D, (3) gives an upper bound on D depending only on δ, λ, and ε. Fix such an upper bound D_0, so that $[p, q] \subseteq \mathcal{N}_{D_0}(c)$.

We will now show that $c \subseteq \mathcal{N}_{R'}([p, q])$, where $R' = D_0(1 + \lambda) + \varepsilon/2$.

Consider a maximal subpath c' of c which lies outside the D_0–neighborhood of $[p, q]$. The closed D–neighborhood of $c|_{[p,c'_-]}$ and of $c|_{[c'_+,q]}$ are both closed sets which collectively cover $[p, q]$. Thus there exists a point w in the intersection of the two closed neighborhoods and $[p, q]$. In other words, there exists $w \in [p, q]$, $t \in c|_{[p,c'_-]}$, and $t' \in c|_{[c'_+,q]}$ such that $d(w, t) \leq D_0$ and $d(w, t') \leq D_0$. In particular, $d(t, t') \leq 2D_0$. Note that c' is a subpath of $c|_{[t,t']}$, and since c is a quasigeodesic, we have $\ell(c') \leq \ell(c|_{[t,t']}) \leq \lambda \cdot 2D_0 + \varepsilon$. Therefore starting from any point on c', we can follow c and reach either t or t' in a distance of at most $\lambda D_0 + \varepsilon/2$. From there we can reach $w \in [p, q]$ in an additional distance of D_0, and so we conclude that $c' \subseteq \mathcal{N}_{R'}([p, q])$. Applying this same argument to all such maximal subpaths, we conclude that $c \subseteq \mathcal{N}_{R'}([p, q])$. $\square$

Definition 2.18. A finitely generated group is *hyperbolic* if there exists a finite generating set X such that $\Gamma(G, S)$ is hyperbolic.

Note that by Corollary 2.16, the "there exists" in the above definition is equivalent to "for all". Moreover, by the Milnor–Schwartz Lemma (Theorem 1.12), an equivalent definition is: A finitely generated group is *hyperbolic* if it admits a proper, cobounded action on a proper, geodesic, hyperbolic metric space.

Examples 2.19. The following groups are hyperbolic.

1. Finite groups
2. $\mathbb{F}_n$ for all $n \geq 1$
3. If M is a closed hyperbolic manifold, then $\pi_1(M)$ is hyperbolic. In particular, if S is an orientable surface of genus g, then $\pi_1(S)$ is hyperbolic if and only if $g \geq 2$.
4. $\mathbb{Z}^n$ is hyperbolic if and only if $n = 1$.

Recall that a finitely generated group is quasi-isometric to any finite-index subgroup. This motivates the following definition of "virtually".

Definition 2.20. Let P be a property of groups. Then we say G is *vitually* P if there is a finite index subgroup $G_0 \leq G$ such that G_0 is P.

For example, finite groups are virtually trivial!

Example 2.21. A virtually hyperbolic group is hyperbolic. For example, the group $PSL_2(\mathbb{Z}) = (\mathbb{Z}/3\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z})$ is virtually free and therefore hyperbolic.

Recall that we showed (as an exercise assigned in class) that $\Gamma(BS(1, 2), \{a, t\})$ is not δ–hyperbolic for any δ. Therefore, the Baumslag-Solitar group $BS(1, 2)$ is *not* a hyperbolic group. Similarly, $BS(1, n)$ is not a hyperbolic group.

We next describe a huge class of examples of hyperbolic groups. To do so, we need to first formally define a group presentation.

Given a set S, we denote by $F(S)$ the free group on S. Elements of $F(S)$ are equivalences classes of words over the alphabet $S^{\pm 1}$: a *word* is a finite sequence $a_1 \dots a_n$ where $a_i \in A^{\pm 1}$. Two words are *equivalent* if you can pass from one to the other by inserting or deleting words of the type ss^{-1}; for example, s_1s_2 is equivalent to $s_1s_1s_1^{-1}s_2$. A words $a_1 \dots a_n$ is *reduced* if $a_i \neq a_{i-1}^{-1}$ for all i. There is a unique reduced word in each equivalence class. If $S = \{s_1, \dots, s_n\}$, we often write $\mathbb{F}_n$ (or $\mathbf{F}_n$ or F_n).

Let G be a group and R a subset of G. The *normal closure* $\langle\langle R \rangle\rangle$ of R in G is the smallest normal subgroup of G containing R. Equivalently,

$$\langle\langle R \rangle\rangle = \left\{ \prod_{i=1}^{k} g_i r_i g_i^{-1} \;\middle|\; k \geq 0,\, g_i \in G,\, r_i \in R^{\pm 1} \right\}.$$

Given a set S and $R \subseteq F(A)$, we say that

$$\langle S \mid R \rangle$$

is a *presentation* for the group $G \simeq F(S)/\langle\langle R \rangle\rangle$. The set S is the set of *generators* of G and R is the set of *relations*. The presentation is *finite* if both S and R are finite sets, and G is said to be *finitely presented* if it admits such a presentation.

Example 2.22. A "random" group is hyperbolic. To be precise: fix $m \geq 2$ and $n \geq 1$. We will temporarily fix $\ell \geq 0$. Consider groups of the form

$$\Gamma = \langle a_1, \dots, a_m \mid r_1, \dots, r_n \rangle,$$

such that r_i are all *cyclically reduced* words of length ℓ (i.e., reduced words such that the first letter is not the inverse of the last letter). Put the uniform probability distribution of the set of all such groups. This defines a group-valued random variable Γ_ℓ. For a property P, we say that a *random group is* P if

$$\lim_{\ell \to \infty} \mathbb{P}(\Gamma_\ell \text{ is } P) = 1$$

for all n, m. This is called the *few-relator model* of random groups.

Theorem 2.23 (Gromov). *A random group is infinite and hyperbolic.*

This is not the only way to define a "random group" (we call the different definitions different *models* of random groups). Generally speaking, any way you define a random group so that you don't always end up with finite groups should yield a hyperbolic group.

A hyperbolic group is *non-elementary* if its Cayley graph with respect to some (equivalently, any) finite generating set is infinite and not quasi-isometric to $\mathbb{R}$. (We call spaces quasi-isometric to $\mathbb{R}$ *quasi-lines*.)