Geometric group theory
Lecture 3

Lawk Mineh

October 2025

1 More hyperbolic geometry

Another key difference between Euclidean and hyperbolic geometry that we highlight here concerns balls. In Euclidean space of dimension n, a ball of radius r has volume proportional to r^n. Already, however, the area of a disc in the hyperbolic plane has grows exponentially with respect to the radius. In fact, the same is true even of the circumference of a circle in hyperbolic space. One may verify this by means of computing some not-too-complicated integrals. As a result, one sees that two different (unit-speed) geodesic rays in $\mathbb{H}^n$ with a shared origin 'diverge' exponentially quickly, in the sense that one must travel exponentially long distances with respect to a radius to get from a point on one to the other, outside a ball around the origin with that radius.

Surfaces and tessellations. In general, it is a fact that every hyperbolic manifold arises as the quotient $\mathbb{H}^n/\Gamma$ of hyperbolic space by a torsion-free discrete subgroup of isometries $\Gamma \leqslant \mathrm{Isom}(\mathbb{H}^n)$. This is easy to see if one assumes that $\mathbb{H}^n$ is the unique simply connected Riemannian manifold of constant curvature -1: each hyperbolic n-manifold M has universal cover $\widetilde{M} = \mathbb{H}^n$ (after possibly rescaling the metric) on which $\Gamma = \pi_1(M)$ acts by isometries. That Γ is torsion-free corresponds to the fact that the action of $\pi_1(M)$ is free and $M = \mathbb{H}^n/\Gamma$ has no singular points. This is a little abstract, so let us restrict our attention now to the two-dimensional case of the hyperbolic plane, where such a realisation can be explicitly computed and visualised.

Our construction will involve understanding polygons and tessellations of the hyperbolic plane. We sketch a proof of the following:

Lemma 1.1. *Let $n, m \in \mathbb{N}$ be natural numbers with $\frac{1}{n} + \frac{1}{m} < \frac{1}{2}$. Then there is a tessellation of $\mathbb{H}^2$ by regular n-gons, with m different n-gons meeting at every vertex.*

Proof. Very close to the origin in the Poincaré disc, the metric closely resembles that of Euclidean space. That is, there are regular n-gons centred on the origin, whose interior angle sum is arbitrarily close to $\frac{1}{2}(n-1)\pi$, the corresponding angle sum in Euclidean space. Following the above discussion on area, moving the vertices outward from the origin decreases this angle sum monotonically, and the sum approaches zero as the n-gon tends to an ideal n-gon. By a continuity argument, there are regular hyperbolic n-gons each of whose interior angles is equal to a given $0 < \theta < \frac{1}{2}(1 - \frac{1}{n})\pi$.

If θ in the above is taken of the form $\frac{2\pi}{m}$ for some natural number m, then we can obtain a tessellation of the hyperbolic plane by regular polygons, by reflecting such a polygon along its edges. The condition that $\theta = \frac{2\pi}{m} < \frac{1}{2}(1 - \frac{1}{n})\pi$ can be rearranged exactly into the hypothesis in the lemma, so it holds by assumption. □

Using this, we may realise hyperbolic manifold structures on every surface that is not the torus or the sphere.

Proposition 1.2. *For each $g \geq 2$, the surface Σ_g of genus g admits a Riemannian metric of constant negative sectional curvature. More precisely, there is a discrete torsion-free subgroup $\Gamma \leqslant \mathrm{Isom}(\mathbb{H}^2)$ such that Σ_g is isometric to $\mathbb{H}^2/\Gamma$.*

Proof. Recall that Σ_g may be realised as a quotient of a (regular) $4g$-gon P. By Lemma 1.1, there is a tessellation of $\mathbb{H}^2$ by copies of P. Now $\mathrm{Isom}(\mathbb{H}^2)$ acts transitively on (oriented) line segments of the same length, so that there are isometries realising each of the side identifications appearing in the above quotient. Let Γ be the subgroup of $\mathrm{Isom}(\mathbb{H}^2)$ generated by these finitely many isometries.

As any element of Γ preserves the given tessellation of $\mathbb{H}^2$, it is straightforward to check that the action is properly discontinuous. Since $\mathbb{H}^2$ is a proper metric space, this means that Γ is a discrete subgroup of $\mathrm{Isom}(\mathbb{H}^2)$. Finally, observe that Γ fixes no points of $\mathbb{H}^2$. The existence of torsion in Γ would imply the existence of a fixed point, so Γ must be torsion-free. □

The above is essentially a simple case of a more general theorem of Poincaré, which constructs discrete subgroups of $\mathrm{Isom}(\mathbb{H}^2)$ whose quotient realises any (orbi)surface with genus g and cone points of order $m_1, \ldots, m_n$, provided

$$2g - 2 + \sum_{i=1}^{n} \left(1 - \frac{1}{m_i}\right) \geq 0.$$

The quantity on the left is often called the *signature* of the surface. One can prove Poincaré's polygon theorem similarly to the above, but with a more involved construction of fundamental domain to account for the cone points.

It follows immediately from the above, together with the Milnor–Schwarz Lemma.

Corollary 1.3. *Let Σ be a surface of genus $g \geq 2$. Then $\pi_1\Sigma$ is quasi-isometric to $\mathbb{H}^2$.*

Discrete subgroups of $\mathrm{Isom}(\mathbb{H}^2)$ are often called *Fuchsian groups*, and discrete subgroups of $\mathrm{Isom}(\mathbb{H}^3)$ – and sometimes those of $\mathrm{Isom}(\mathbb{H}^n)$ – are called *Kleinian*.

The boundary. Any element of $\mathrm{Isom}(\mathbb{H}^n)$ has an induced action by homeomorphisms on the boundary $\partial\mathbb{H}^n$ of hyperbolic space, so there is a well-defined homomorphism $\mathrm{Isom}(\mathbb{H}^n) \to \mathrm{Homeo}(S^{n-1})$ for each $n \in \mathbb{N}$. As such, we can attempt to retrieve information about subgroups of isometries of $\mathbb{H}^n$ by analysing their action on the boundary. For individual isometries, this turns out to be very doable.

Proposition 1.4. *Let $g \in \mathrm{Isom}(\mathbb{H}^n)$ be an isometry. Then either*

1. *g fixes a point in $\mathbb{H}^n$;*
2. *g fixes exactly one point in $\partial\mathbb{H}^n$; or*
3. *g fixes exactly two points in $\partial\mathbb{H}^n$,*

and these possibilities are mutually exclusive.

We call the isometries *elliptic*, *parabolic*, and *loxodromic* respectively in the above cases. In each of the above cases, the geometry of the isometry g can be effectively described. If g is elliptic, then it is a rotation around its fixed point in $\mathbb{H}^n$, since the stabiliser of any point in $\mathbb{H}^n$ is the Lie group $O(n)$. If g is parabolic, then it fixes any *horosphere* centred around its fixed point in $\partial\mathbb{H}^n$. A horosphere is the limit of a sequence of spheres of increasing radii with a shared point of tangency, and its centre is the point in the boundary that meets the diameter of these spheres. A horosphere is a copy of the Euclidean plane, embedded in $\mathbb{H}^n$ with exponential distortion of the metric. Lastly, a loxodromic isometry fixes a bi-infinite geodesic joining its two fixed points in $\partial\mathbb{H}^n$, and acts as a translation when restricted to this axis.

For more general subgroups of isometries than cyclic ones, the situation is naturally more complicated. Here, the action of the subgroup on a particular subset of the boundary known as its *limit set* becomes important.

Definition 1.5. Let $G \leqslant \mathrm{Isom}(\mathbb{H}^n)$. The *limit set* of G is the subset $\Lambda G \subseteq \partial\mathbb{H}^n$ of accumulation points of G-orbits in $\mathbb{H}^n$.

Exercise 1.6. Show that ΛG is the smallest G-invariant closed subset of $\partial\mathbb{H}^n$, and that ΛG is a perfect compactum unless G has a finite index cyclic subgroup.

The study of limit sets and their geometric properties is of much interest. They are in general fractal subsets of the boundary. Various facts about a subgroup can be determined from its limit set; we do not pursue these here, but will return to the topic when discussing boundaries of abstract hyperbolic groups in a later section.

2 Hyperbolic metric spaces

We now introduce a notion of negative curvature for metric spaces. Our definition will be modelled on a key property of the classical hyperbolic spaces of the previous subsection: it will state that every geodesic triangle is uniformly thin. That is, triangles in these spaces will look somewhat like tripods. Triangles are the most basic shapes in a geodesic space, and so, as we will see, this assumption has some strong consequences for the geometry of these spaces and the groups that act on them.

We will need some preliminary definitions.

Definition 2.1 (Gromov product)**.** Let (X, d) be a metric space, and $x, y, z \in X$ be points. The *Gromov product* of x and y with respect to z is

$$\langle x, y \rangle_z = \frac{1}{2}\Big(\mathrm{d}(x, z) + \mathrm{d}(y, z) - \mathrm{d}(x, y)\Big).$$

One can think of the Gromov product $\langle x, y\rangle_z$ as an abstracted notion of the 'angle' spanned by x and y with respect to z. Indeed, in Euclidean space, this Gromov product is exactly the distance of the point z to the points on $[x, z]$ and $[y, z]$ that touch the incircle of the triangle with vertices x, y, and z – up to homothety, this is determined by the angle these two lines make.

Definition 2.2 (Thin triangles)**.** Let Δ be a geodesic triangle with vertices x, y, and z in a metric space X, and let $\delta \geq 0$. Call T_Δ the tripod with leg lengths $\langle x, y\rangle_z, \langle x, z\rangle_y$, and $\langle y, z\rangle_z$. There is a unique map $\varphi\colon \Delta \to T_\Delta$ such that x, y, and z map to the extremal vertices of T_Δ and φ restricts to an isometry on each side of Δ. We say Δ is *δ-thin* if $\operatorname{diam} \varphi^{-1}(\{t\}) \leq \delta$ for all $t \in T_\Delta$.

Definition 2.3 (Hyperbolic metric space)**.** Let X be a geodesic metric space. If there is $\delta \geq 0$ such that every geodesic triangle in X is δ-thin, we say that X is a *δ-hyperbolic metric space*. We simply call X a *hyperbolic metric space* if there is some $\delta \geq 0$ such that it is a δ-hyperbolic metric space.

Exercise 2.4. Show that a geodesic space is 0-hyperbolic if and only if it is an $\mathbb{R}$-*tree*: a space in which every pair of points is connected by a unique arc.

Example 2.5. We saw in the previous subsection that $\mathbb{H}^n$ is δ-hyperbolic with hyperbolicity constant $\delta = \frac{1}{2}\log 3$.

Example 2.6. The plane $\mathbb{R}^2$ is not a hyperbolic metric space, as for any $\delta \geq 0$, any equilateral triangle with side lengths greater than 2δ is not δ-thin.

There are many alternative formulations of the thin triangles condition. One that is very commonly used and can be useful is the *slim triangles* formulation.

Definition 2.7. Let Δ be a geodesic triangle in metric space X, and let $\delta \geq 0$. We say that Δ is *δ-slim* if each side of Δ is contained in a δ-neighbourhood of the union of the other two sides.

Of course, using slim triangles instead of thin triangles gives an identical characterisation of hyperbolic metric spaces, up to a small change in the constant in the definitions.

Exercise 2.8. Show that every δ-thin triangle is δ-slim, and that every δ-slim triangle is also 2δ-thin.

Exercise 2.9. Show that a if geodesic space X is δ-hyperbolic then it satisfies the *four-point condition*: for all $x, y, z, w \in X$, we have

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

Further, show that if X satisfies the four-point condition for some $\delta \geq 0$, there is $\delta' \geq 0$ such that X is δ'-hyperbolic.

Remark 2.10. The four-point condition above can of course be formulated for any metric space, without any assumption on whether or not geodesics exist. This is frequently useful, as it allows us to talk about hyperbolicity of discrete metric spaces, such as groups equipped with a word metric.

An important feature of hyperbolic metric spaces is the 'stability' of quasi-geodesics: all quasi-geodesics actually follow geodesics between their endpoints uniformly closely. This fact is usually referred to as the *Morse Lemma*. Note that this feature is particular to hyperbolic metric spaces; the following example shows that it can dramatically fail outside of this setting.

Example 2.11. Consider the Cayley graph of $\mathbb{Z}^2$ with the standard generating set. Then take the concatenation of three geodesics of length n, one vertical geodesic going up, one horizontal going right, and one vertical going down, is a $(3,0)$-quasi-geodesic. This path contains points that are a distance of n from the unique geodesic joining its endpoints (which is a horizontal path of length n), and we may take n to be arbitrarily large.

In order to obtain the stability statement, we must first obtain an estimate on the length of paths different from geodesics. Simply stated, we have that the length of a path grows at least exponentially with its distance from a geodesic between its endpoints.

Lemma 2.12. *Let X be a δ-hyperbolic space, and $x, y \in X$. If $\gamma\colon I \to X$ is a continuous rectifiable path between x and y, then*

$$\mathrm{d}(z, \gamma(I)) \leq \delta \max\{0, \log_2 \ell(\gamma)\} + 2$$

for any z on a geodesic between x and y.

Proof. Let $n = \lceil \log_2 \ell(\gamma) \rceil$ and suppose that γ is a parameterisation proportional to arc-length, so that we may write $I = [0,1]$. If $\ell(\gamma) \leq 1$ then there is nothing to prove, so suppose otherwise. It follows that $n \geq 1$ is a positive natural number.

Let z be a point on a geodesic from $x = \gamma(0)$ to $y = \gamma(1)$. Since triangles in X are δ-slim, for any $k \geq 0$ and $i = 1, \ldots, 2^k$, any point on a geodesic $[\gamma(\frac{i-1}{2^k}), \gamma(\frac{i}{2^k})]$ is at most distance δ from a geodesic of the form $[\gamma(\frac{j-1}{2^{k+1}}), \gamma(\frac{j}{2^{k+1}})]$ for some $1 \leq j \leq 2^{k+1}$. It follows by a finite induction that for any $k \geq 0$, there is $1 \leq i \leq 2^k$ such that

$$\mathrm{d}(z, p) \leq \delta k. \tag{2.1}$$

where p is a geodesic of the form $[\gamma(\frac{i-1}{2^k}), \gamma(\frac{i}{2^k})]$.

Now for any $i = 1, \ldots, 2^n$, the length of the subpath γ restricted to $[\frac{i-1}{2^n}, \frac{i}{2^n}]$ is at most 1 by the choice of n. Hence any point on a geodesic between $\gamma(\frac{i-1}{2^n})$ and $\gamma(\frac{i}{2^n})$ is a distance of less than 1 from $\gamma(I)$. Combined with (2.1) applied to the case $k = n$, this fact gives the required inequality. ☐

Proposition 2.13 (Morse Lemma)**.** *Let X be a δ-hyperbolic space, $\lambda \geq 1$, and $c \geq 0$. There is a constant $M = M(\lambda, c, \delta) \geq 0$ such that any (λ, c)-quasi-geodesic segment in X is a Hausdorff distance of at most M from any geodesic segment between its endpoints.*

Proof (first half): Let $\gamma\colon I \to X$ be a (λ, c)-quasi-geodesic segment, and write x and y for its endpoints. Applying the lemma on continuous quasi-geodesics, there is a (λ, c')-quasi-geodesic $\gamma'\colon J \to X$ with the same endpoints as γ, lying a Hausdorff distance of at most c' from γ. We may suppose that γ' is parametrised by arc-length. Fix a geodesic p whose endpoints are also x and y. Let z be a point on p maximising $r = \mathrm{d}(z, \gamma')$, which exists by continuity. We will exhibit an absolute bound on r.

Let x' and y' be points on p between x and z and y and z respectively, with $\mathrm{d}(x', z) = \mathrm{d}(y', z) = 2r$ (choosing $x' = x$ or $y' = y$ if $\mathrm{d}(x, z) \leq 2r$ or $\mathrm{d}(y, z) \leq 2r$ respectively). By the definition of z, we have

$$\mathrm{d}(x', \gamma') \leq r \quad \text{and} \quad \mathrm{d}(y', \gamma') \leq r.$$

Let $s, t \in [a, b]$ be points such that $\mathrm{d}(x', \gamma'(s))$ and $\mathrm{d}(y', \gamma'(t))$ realise the distances from x' and y' to γ', which again exist by continuity. We will say $s \leq t$, swapping the names if otherwise. It follows then, that

$$\mathrm{d}(\gamma'(s), \gamma'(t)) \leq \mathrm{d}(\gamma'(s), x') + \mathrm{d}(x', y') + \mathrm{d}(y', \gamma'(t)) \leq 6r$$

Write ξ for concatenation of the subpath of γ' between $\gamma'(s)$ and $\gamma'(t)$ with geodesics $[x', \gamma'(s)]$ and $[\gamma'(t), y']$. Now using the fact that γ' is a quasi-geodesic, the above implies that

$$\ell(\xi) \leq r + \ell(\gamma'|_{[s,t]}) + r \leq (6\lambda + 2)r + \lambda c'.$$

Then by Lemma 2.12 and the choice of z, we have

$$r = \mathrm{d}(z, \xi) \leq \delta \log_2((6\lambda + 2)r + \lambda c') + 2.$$

Rearranging slightly, we have

$$2^{\frac{1}{\delta}(r-1)} \leq (6\lambda + 2)r + \lambda c'$$

so that an exponential function in r is bounded above by a linear function in r. This imposes a uniform bound M' on r depending on these two functions, which in turn depend only on δ, λ, and c. Thus p is contained in a M'-neighbourhood of γ'. $\square$