Introduction to Riemannian manifolds

All manifolds will be connected, Hausdorff and second countable.

Terminology. Let M be a smooth manifold. Denote the tangent space at $x \in M$ by T_xM. If $f\colon M \to N$ is a smooth map between smooth manifolds, denote the associated map on T_xM by $(Df)_x\colon T_xM \to T_{f(x)}N$. If I is an open interval in $\mathbb{R}$ and $\alpha\colon I \to M$ is a smooth path, then for $t \in I$, $\alpha'(t)$ denotes $(D\alpha)_t(1) \in T_{\alpha(t)}M$.

Definition. A *Riemannian metric* on a smooth manifold M is a choice at each point $x \in M$ of a positive definite inner product $\langle\ ,\ \rangle$ on T_xM, the inner products varying smoothly with x. Then M is known as a *Riemannian manifold*. We will not give a formal definition of the phrase 'varying smoothly with x'.

Definition. A *local isometry* between two Riemannian manifolds M and N is a local diffeomorphism $h\colon M \to N$, such that, for all points $x \in M$ and all vectors v and w in T_xM,

$$\langle v, w \rangle = \langle (Dh)_x(v), (Dh)_x(w) \rangle.$$

A *(Riemannian) isometry* is a local isometry that is also a diffeomorphism.

Let M be a Riemannian manifold and let x be a point in M. The Riemannian metric allows one to define for a vector $v \in T_xM$ the length $||v|| = \langle v, v\rangle^{1/2}$ and also the angle between two non-zero vectors v and w in T_xM:

$$\cos(\text{Angle}(v, w)) = \frac{\langle v, w \rangle}{||v||\,||w||}.$$

The lengths $||\ ||$ determine the inner product: if $v, w \in T_xM$, then

$$\langle v, w \rangle = (||v + w||^2 - ||v||^2 - ||w||^2)/2.$$

So, a diffeomorphism which preserves the lengths of vectors is necessarily a Riemannian isometry.

Smooth paths $\alpha\colon [0, T] \to M$ inherit a length, given by

$$\text{Length}(\alpha) = \int_0^T ||\alpha'(t)||\, dt.$$

This is independent of its parametrisation - in other words, if $\beta\colon [0, T_1] \to [0, T]$ is a diffeomorphism, then $\text{Length}(\alpha \circ \beta) = \text{Length}(\alpha)$. This is just a consequence of the fact that we can change the variable in the integration. A *piecewise smooth* path $\alpha\colon [0,T] \to M$ is a path that is smooth at all but finitely many points. Piecewise smooth paths also inherit a length. We construct a metric d on M: if x and y are points in M, then

$$d(x,y) = \inf\{\text{Length}(\alpha) : \alpha \text{ is a piecewise smooth path from } x \text{ to } y\}.$$

Proposition 1.1. *This does give a metric on M. The topology induced by this metric coincides with the original topology on M.*

Notation. If x is a point in a metric space M and $\epsilon > 0$, denote $\{y \in M : d(x,y) < \epsilon\}$ by $B_\epsilon(x)$.

Crucial in the study of Riemannian manifolds is the notion of a geodesic. Here's a non-standard definition, which is equivalent to the usual one.

Definition. A *geodesic* (with speed $s \in \mathbb{R}_{\geq 0}$) is a smooth map $\alpha\colon I \to M$ (where I is an interval in $\mathbb{R}$) such that $||\alpha'(t)|| = s$ for all $t \in I$ and which is 'locally length minimising'. This means that for all $t \in I$, there is an $\epsilon > 0$, such that for all t_1 and t_2 in $(t-\epsilon, t+\epsilon) \cap I$,

$$d(\alpha(t_1), \alpha(t_2)) = s|t_1 - t_2|.$$

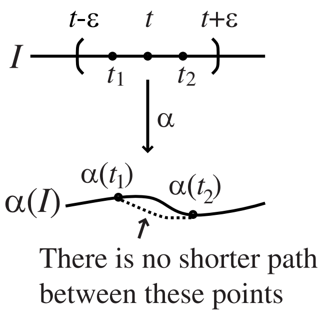

Figure 1.

Exercise. 1. The geodesics in $\mathbb{R}^n$ are straight lines.

2. The geodesics in S^2 are great circles.

Remarks. 1. This demonstrates that geodesics need not be globally length-minimising. In other words, it need not be true that

$$d(\alpha(t_1), \alpha(t_2)) = s|t_1 - t_2|$$

for all $t_1, t_2 \in I$. For example, great circles in S^2. This example also demonstrates that there need not be a unique shortest path between two points.

2. The maximal interval $I \subset \mathbb{R}$ on which a geodesic is defined need not be the whole of $\mathbb{R}$. For example, consider geodesics in the open unit disc in $\mathbb{R}^2$.

3. There need not be a shortest path between two points. For example, consider the points $(-1, 0)$ and $(1, 0)$ in $\mathbb{R}^2 - \{(0, 0)\}$. But if there is a shortest path between two points, then we may find one which has constant speed. This is necessarily globally length minimising and hence a geodesic.

4. A local isometry between Riemannian manifolds (for example, the inclusion of an open subset) preserves geodesics.

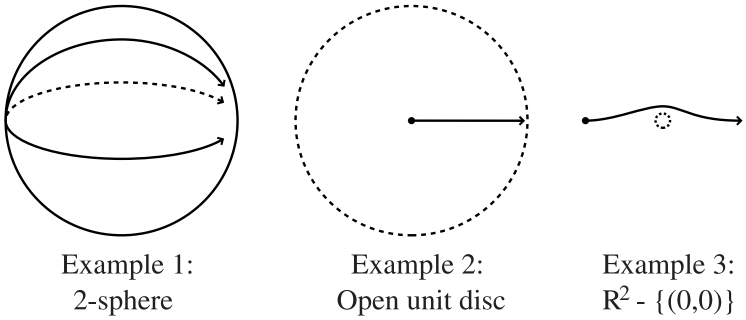

Figure 2.

A fundamental result from differential geometry is the following.

Theorem 1.2. [Existence and uniqueness of geodesics] *For all points $x \in M$ and for all $v \in T_xM$, there is a unique maximal interval $I \subset \mathbb{R}$ containing a neighbourhood of* 0, *and a unique geodesic $\alpha: I \to M$, such that $\alpha(0) = x$ and $\alpha'(0) = v$.*

Idea of proof. Pick a chart $\phi: U \to \mathbb{R}^n$ around x. For each path $\alpha: [-T, T] \to U$, consider $\phi \circ \alpha: [-T, T] \to \mathbb{R}^n$. One shows that α is a geodesic if and only if the n co-ordinates of $\phi \circ \alpha$ satisfy certain second order differential equations. These differential equations have a solution for small enough T, which is unique given the initial conditions $\alpha(0) = x$ and $\alpha'(0) = v$. □

Definition. The *exponential map* at a point $x \in M$ is the map $\exp_x$ from a subset of T_xM to a subset of M which takes a vector $v \in T_xM$ to $\alpha(1)$, where $\alpha: I \to M$

is the geodesic from Theorem 1.2 with $\alpha(0) = x$ and $\alpha'(0) = v$, *providing* $1 \in I$.

Proposition 1.3. *For each point* $x \in M$, $\exp_x$ *is a smooth map, whose domain is an open neighbourhood of* 0. *For small enough* $\epsilon > 0$, $\exp_x$ *maps* $B_\epsilon(0) \subset T_xM$ *diffeomorphically onto* $B_\epsilon(x) \subset M$.

Idea of proof. As in Theorem 1.2, one relates geodesics to certain second order differential equations, and then one uses the fact that their solutions are smooth and depend smoothly on the initial conditions. For the second part, one first determines the derivative of $\exp_{\mathrm{x}}$ and discovers that it has maximal rank. Hence, the inverse function theorem gives that $\exp_x$ sends $B_\epsilon(0) \subset T_xM$ diffeomorphically onto its image in M, which is clearly $B_\epsilon(x) \subset M$. □

Proposition 1.4. *If* $h\colon M \to N$ *is a local isometry between Riemannian manifolds, and* $x \in M$, *then the following diagram commutes (where the maps are defined):*

$$\begin{array}{ccc} T_xM & \xrightarrow{(Dh)_x} & T_{h(x)}N \\ \Big\downarrow \exp_x & & \Big\downarrow \exp_{h(x)} \\ M & \xrightarrow{h} & N \end{array}$$

Proof. Pick $v \in T_xM$. Let α be the unique geodesic in M with $\alpha(0) = x$ and $\alpha'(0) = v$. Since h is a local isometry, it preserves geodesics and so $h \circ \alpha$ is a geodesic in N. But $(h \circ \alpha)(0) = h(x)$ and $(h \circ \alpha)'(0) = (Dh)_x(v)$. Therefore, the uniqueness part of Theorem 1.2 gives that $h(\exp_x(v)) = (h \circ \alpha)(1) = \exp_{h(x)}((Dh)_x(v))$. □

Theorem 1.5. *Let* M *and* N *be Riemannian manifolds, with* M *connected. Let* $h\colon M \to N$ *and* $k\colon M \to N$ *be local isometries onto their images. Suppose that for some* $x \in M$, $h(x) = k(x)$ *and* $(Dh)_x = (Dk)_x$. *Then* $h = k$.

Proof. Consider the set

$$U = \{y \in M : h(y) = k(y) \text{ and } (Dh)_y = (Dk)_y\colon T_yM \to T_{h(y)}N\}.$$

We first show that U is open. Pick $y \in U$. By Proposition 1.4, the following diagram commutes:

$$\begin{array}{ccccc} T_{h(y)}M & \xleftarrow{(Dh)_y} & T_yM & \xrightarrow{(Dk)_y} & T_{k(y)}M \\ \Big\downarrow \exp_{h(y)} & & \Big\downarrow \exp_y & & \Big\downarrow \exp_{k(y)} \\ N & \xleftarrow{h} & M & \xrightarrow{k} & N \end{array}$$

But $h(y) = k(y)$ and $(Dk)_y = (Dh)_y$. Therefore $h = k$ on the image of $\exp_y$, which is a neighbourhood of y by Proposition 1.3. If $h = k$ on an open set, then $(Dh) = (Dk)$ there. Therefore, U is open.

Now, we show that U is closed. Let $\{y_i : i \in \mathbb{N}\}$ be a sequence of points in U, tending to some point y in M. Then $h(y) = \lim_{i\to\infty} h(y_i) = \lim_{i\to\infty} k(y_i) = k(y)$. Similarly, $(Dh)_y = (Dk)_y$. So, $y \in U$. Therefore, U is closed. Since U is open, closed and non-empty, and M is connected, $U = M$. Therefore $h = k$. □

Remark. A Riemannian manifold M has a (possibly infinite) volume. For each $x \in M$, the paralleliped in T_xM spanned by n orthonormal vectors is defined to have volume 1. By integrating over M, this determines its volume. Compact Riemannian manifolds always have finite volume.

Hyperbolic Manifolds

Hilary Term 2000

Marc Lackenby

Geometry and topology is, more often than not, the study of manifolds. These manifolds come in a variety of different flavours: smooth manifolds, topological manifolds, and so on, and many will have extra structure, like complex manifolds or symplectic manifolds. All of these concepts can be brought together into one overall definition.

A *pseudogroup* on a (topological) manifold X is a set $\mathcal{G}$ of homeomorphisms between open subsets of X satisfying the following conditions:

1. The domains of the elements of $\mathcal{G}$ must cover X.
2. The restriction of any element of $\mathcal{G}$ to any open set in its domain is also in $\mathcal{G}$.
3. The composition of two elements of $\mathcal{G}$, when defined, is also in $\mathcal{G}$.
4. The inverse of an element of $\mathcal{G}$ is in $\mathcal{G}$.
5. The property of being in $\mathcal{G}$ is 'local', that is, if $g: U \to V$ is a homeomorphism between open sets of X, and U has a cover by open sets U_α such that $g|_{U_\alpha}$ is in $\mathcal{G}$ for each U_α, then g is in $\mathcal{G}$.

For example, the set of all diffeomorphisms between open sets of $\mathbb{R}^n$ forms a pseudogroup.

A *$\mathcal{G}$-manifold* is a Hausdorff topological space M with a countable $\mathcal{G}$-atlas. A *$\mathcal{G}$-atlas* is a collection of $\mathcal{G}$-compatible co-ordinate charts whose domains cover M. A *co-ordinate chart* is a pair (U_i, ϕ_i), where U_i is an open set in M and $\phi_i: U_i \to X$ is a homeomorphism onto its image. That these are *$\mathcal{G}$-compatible* means that whenever (U_i, ϕ_i) and (U_j, ϕ_j) intersect, the transition map $\phi_i \circ \phi_j^{-1}: \phi_j(U_i \cap U_j) \to \phi_i(U_i \cap U_j)$ is in the pseudogroup $\mathcal{G}$.

Unless otherwise stated, all manifolds we will consider will be connected, Hausdorff and second countable.

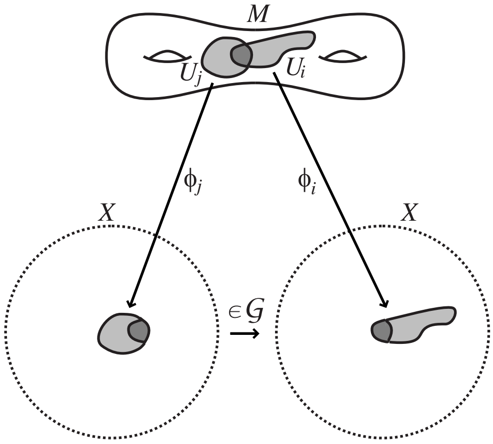

Figure 1.

Examples.

X	Pseudogroup $\mathcal{G}$	$\mathcal{G}$-manifold
$\mathbb{R}^n$	All homeomorphisms between open subsets of $\mathbb{R}^n$	Topological manifold
$\mathbb{R}^n$	All C^r-diffeomorphisms between open subsets of $\mathbb{R}^n$ $(1 \leq r \leq \infty)$	Differentiable manifold (of class C^r)
$\mathbb{C}^n$	All biholomorphic maps between open subsets of $\mathbb{C}^n$	Complex manifold

Other examples. Real analytic manifolds, foliated manifolds, contact manifolds, symplectic manifolds, piecewise linear manifolds.

The above definition of a $\mathcal{G}$-manifold was actually a little ambiguous. When is it possible for two different $\mathcal{G}$-atlases to define the same $\mathcal{G}$-structure? Two $\mathcal{G}$-atlases on a topological space M define the same $\mathcal{G}$-structure if they are *compatible*, which means that their union is also a $\mathcal{G}$-atlas. Compatibility is an equivalence relation (exercise) and hence any $\mathcal{G}$-atlas is contained in a well-defined equivalence class of $\mathcal{G}$-atlases.

Exercise. Let $\mathcal{G}$ be the set of translations of $\mathbb{R}$ restricted to open subsets of $\mathbb{R}$. Show that $\mathcal{G}$ satisfies the first four conditions in the definition of a pseudogroup, but fails the fifth condition. Show that compatibility between $\mathcal{G}$-atlases on S^1 is not an equivalence relation.

Let M be a $\mathcal{G}$-manifold and let $h: N \to M$ be a local homeomorphism (that is, each point of N has an open neighbourhood U such that $h|_U$ is an open mapping that is a homeomorphism onto its image). Then we may pull back the $\mathcal{G}$-structure on M to a $\mathcal{G}$-structure on N.

A homeomorphism $h: N \to M$ between $\mathcal{G}$-manifolds is a *$\mathcal{G}$-isomorphism* if the pull back $\mathcal{G}$-structure on N is the same as the $\mathcal{G}$-structure it possesses already.

Let $\mathcal{G}_0$ be a collection of homeomorphisms between open subsets of a manifold X. The pseudogroup $\mathcal{G}$ *generated by* $\mathcal{G}_0$ is the intersection of all pseudogroups on X containing $\mathcal{G}_0$. It is the smallest pseudogroup containing $\mathcal{G}_0$.

In certain cases, it is possible to identify the pseudogroup that is generated much more explicitly.

Special case. Let G be a group acting on a manifold X. Let $\mathcal{G}$ be the pseudogroup generated by G. Then $g \in \mathcal{G}$ if and only if the domain of g can be covered by open sets U_α such that $g|_{U_\alpha} = g_\alpha|_{U_\alpha}$ for some $g_\alpha \in G$ (exercise). A $\mathcal{G}$-manifold is also called a (G, X)-*manifold*.

Terminology.

X	G	(G, X)-manifold
$\mathbb{R}^n$	Euclidean isometries	Euclidean manifold
S^n	Spherical isometries	Spherical manifold
$\mathbb{R}^n$	Affine transformations	Affine manifold
$\mathbb{R}^n$	Euclidean similarities	Similarity manifold

In each of these cases, the group G is quite small (much smaller than the full diffeomorphism pseudogroup) and so the resulting (G, X)-structures are quite rigid.

Examples. 1. By taking a single chart, any open subset of $\mathbb{R}^n$ is a (G, X)-manifold for all (G, X).

2. The torus admits a Euclidean structure, with the following charts.

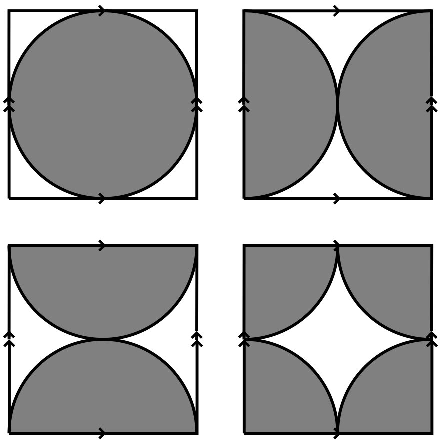

Figure 2.

Another way of constructing this example is as follows. Let M be a manifold and let $\tilde{M}$ be its universal cover, with G the group of covering transformations. Then M inherits a $(G, \tilde{M})$-structure.

The action of a group G on a manifold X is *rigid* if, whenever two elements of G agree on an open set of X, they are the same element of G. Then the pseudogroup generated by such a G is the set of homeomorphisms $h\colon U \to h(U)$ between open subsets of X such that the restriction of h to any component of U is the restriction of an element of G. The examples above of groups G acting on a manifold X are all rigid. Also, if X is a Riemannian manifold and G is a group of isometries of X, then G acts rigidly. (This is a consequence of Theorem 1.5 of the Introduction to Riemannian Manifolds.)

Euclidean structures are very well understood, as demonstrated by the following result.

Theorem. [Bieberbach] *Every closed Euclidean n-manifold is finitely covered by a torus T^n.*

For example, the only closed surfaces that support Euclidean structures are the torus and the Klein bottle. Spherical structures are even more restrictive.

Theorem. *A closed spherical n-manifold is finitely covered by S^n. In particular, it has finite fundamental group.*

There is a fascinating conjectured converse to this in dimension three.

Conjecture. *Every closed 3-manifold with finite fundamental group admits a spherical structure.*

This implies the famous Poincaré conjecture.

Poincaré Conjecture. *A closed 3-manifold with trivial fundamental group is homeomorphic to S^3.*

In this course we will define and study another model space X. We will define, for each $n \geq 1$, a Riemannian n-manifold $\mathbb{H}^n$, known as *hyperbolic space*. Its isometry group is denoted by $\text{Isom}(\mathbb{H}^n)$. An $(\text{Isom}(\mathbb{H}^n), \mathbb{H}^n)$-manifold is known as a *hyperbolic manifold*. A hyperbolic manifold inherits a Riemannian metric.

It is a theorem from Riemannian geometry that $\mathbb{H}^n$ (respectively, S^n, Euclidean space) is the unique complete simply-connected Riemannian n-manifold with all sectional curvatures being -1 (respectively, one, zero). Hyperbolic manifolds are precisely those Riemannian manifolds in which all sectional curvatures are -1.

Hyperbolic space has a richer isometry group than Euclidean or spherical space, and hence it will be easier to find hyperbolic structures. But still, hyperbolic manifolds are sufficiently rigid to have interesting properties. Here are some sample results about hyperbolic manifolds.

A smooth 3-manifold is *irreducible* if any smoothly embedded 2-sphere bounds a 3-ball. A smooth 3-manifold M is *atoroidal* if any $\mathbb{Z} \oplus \mathbb{Z}$ subgroups of $\pi_1(M)$ is conjugate to $i_*(\pi_1(X))$, where $i\colon X \to M$ is the inclusion of a toral boundary component of M. A compact orientable 3-manifold M is *Haken* if it is irreducible and it contains a compact orientable embedded surface S (other than a 2-sphere) with $\partial S = S \cap \partial M$, such that the map $\pi_1(S) \to \pi_1(M)$ induced by inclusion is an injection. Haken 3-manifolds form a large class. In particular, any compact orientable irreducible 3-manifold M with non-empty boundary or with infinite $H_1(M)$ is Haken.

Theorem. [Thurston] *Let M be a closed atoroidal Haken 3-manifold. Then M admits a hyperbolic structure.*

This is a special case of the so-called geometrisation conjecture.

Geometrisation Conjecture. [Thurston] *Any closed irreducible atoroidal 3-manifold admits either a hyperbolic structure or a spherical structure.*

The closed irreducible toroidal 3-manifolds with $\mathbb{Z} \oplus \mathbb{Z}$ subgroups in their fundamental group are known to admit a certain type of 'geometric structure', but the spaces X on which they are modelled have slightly less natural geometries.

The above theorems and conjectures suggest that it may be rather too easy to put a hyperbolic structure structure on a manifold. But in fact this is not the case.

Theorem. [Mostow Rigidity] *Let M and N be closed hyperbolic n-manifolds, with $n > 2$. If $\pi_1(M)$ and $\pi_1(N)$ are isomorphic, then M and N are isomorphic hyperbolic manifolds.*

This is very strong indeed. It says that each of the following implications can be reversed for closed hyperbolic n-manifolds for $n > 2$:

$$\text{Isomorphic} \Rightarrow \text{Isometric} \Rightarrow \text{Diffeomorphic}$$

$$\Rightarrow \text{Homeomorphic} \Rightarrow \text{Homotopy equivalent} \Rightarrow \text{Isomorphic } \pi_1$$

There are lots of geometric invariants of hyperbolic manifolds. For example, they have a well-defined volume. Thus Mostow Rigidity implies that these depend only on the fundamental group of the manifold. In particular, they are topological invariants.

In the case of hyperbolic manifolds, it is those that are complete which are particularly interesting. Mostow's rigidity theorem remains true when the word 'closed' is replaced by 'complete and finite volume'.

Thurston's theorem on the hyperbolisation of closed atoroidal Haken manifolds extends the bounded case as follows.

Theorem. [Thurston] *Let M be a compact orientable irreducible atoroidal 3-manifold, such that ∂M is a non-empty collection of tori. Then either $M - \partial M$ has a complete finite volume hyperbolic structure, or M is homeomorphic to one of the following exceptional cases:*

1. $S^1 \times [0,1] \times [0,1]$

2. $S^1 \times S^1 \times [0,1]$

3. the space obtained by gluing the faces of a cube as follows: arrange the six faces into three opposing pairs; glue one pair, by translating one face onto the other; glue another pair, by translating one face onto the other and then rotating through π *about the axis between the two faces. (This is the total space of the unique orientable* I*-bundle over the Klein bottle.)*

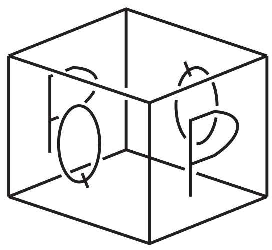

Figure 3.

Example. Let K be a knot in S^3, that is, a smoothly embedded simple closed curve. Let $N(K)$ be an open tubular neighbourhood of K. Then $M = S^3 - N(K)$ is a 3-manifold with boundary a torus, which is compact, orientable and irreducible. Irreducibility is a consequence of the Schoenflies theorem. Note that $M - \partial M$ is homeomorphic to $S^3 - K$. The knots K for which M fails to be atoroidal fall into one of two classes:

1. *torus knots*, which are those that lie on the boundary of a standardly embedded solid torus in S^3, and are not the unknot, and
2. *satellite knots*, which are those that have an embedded π_1-injective torus in their complement that is not boundary-parallel. Such a torus bounds a 'knotted' solid torus in S^3 containing the knot.

So, providing K is neither the unknot, a torus knot nor satellite knot, $S^3 - K$ admits a complete, finite volume hyperbolic structure.

1. Three models for hyperbolic space

Hyperbolic space was discovered by a number of people, including Bolyai, Gauss and Lobachevsky. It has many of the properties of Euclidean space, including:

1. Between any two points in $\mathbb{H}^n$, there is a unique geodesic.

2. For any two points $x, y \in \mathbb{H}^n$, there is an isometry taking x to y.

   However, there is a major difference between $\mathbb{H}^2$ and $\mathbb{R}^2$:

3. If α is a geodesic in $\mathbb{H}^2$ and x is a point not on α, then there are infinitely many geodesics through x which do not meet α.

Remark. We will often confuse a geodesic $\alpha\colon [0,T] \to M$ with its image in M. Thus 'unique geodesic' really means 'unique up to re-parametrisation'.

There are three main 'models' for hyperbolic space, each of which is a Riemannian manifold, any two of which are isometric. Each will be denoted by $\mathbb{H}^n$.

The Poincaré disc model

For each $n \in \mathbb{N}$, let D^n be the open unit ball $\{x \in \mathbb{R}^n\colon d_{Eucl}(x, 0) < 1\}$, where $d_{Eucl}(x, 0)$ is the Euclidean distance between x and the origin 0 in $\mathbb{R}^n$. Assign a Riemannian metric to D^n by defining the inner product of two vectors v and w in $T_x D^n$ to be

$$\langle v, w\rangle_{D^n} = \langle v, w\rangle_{Eucl} \left(\frac{2}{1 - [d_{Eucl}(x,0)]^2}\right)^2,$$

where $\langle\ ,\ \rangle_{Eucl}$ is the standard Euclidean inner product. This is the Poincaré disc model for $\mathbb{H}^n$.

Remarks. 1. Since $\langle v, w\rangle_{D^n}$ is a multiple of $\langle v, w\rangle_{Eucl}$, the angle between two non-zero vectors in $T_x\mathbb{H}^n$ is just their Euclidean angle.

2. The factor $2/(1 - [d_{Eucl}(x,0)]^2) \to \infty$ as $d_{Eucl}(x,0) \to 1$, so hyperbolic distances get very big as $d_{Eucl}(x,0) \to 1$.

3. The inclusions $D^1 \subset D^2 \subset D^3 \subset ...$ induce inclusions $\mathbb{H}^1 \subset \mathbb{H}^2 \subset \mathbb{H}^3$ which respect the Riemannian metrics.

The unit sphere $\{x \in \mathbb{R}^n : d_{Eucl}(x, 0) = 1\}$ is known as the *sphere at infinity* S^{n-1}_∞. It is *not* part of hyperbolic space. But it is nonetheless useful when studying $\mathbb{H}^n$.

The upper half-space model

This is another way of describing hyperbolic space. Let $U^n = \{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_n > 0\}$. Give it a Riemannian metric by defining the inner product of v and w in $T_{(x_1,\ldots,x_n)}U^n$ to be

$$\langle v, w\rangle_{U^n} = \frac{\langle v, w\rangle_{Eucl}}{x_n^2}.$$

Proposition 1.1. *There is a Riemannian isometry between D^n and U^n.*

Proof. Let $\pm e_n = (0, 0, \ldots, \pm 1) \in \mathbb{R}^n$. Consider the map

$$\mathbb{R}^n - \{-e_n\} \xrightarrow{I} \mathbb{R}^n - \{-e_n\}$$
$$x \mapsto 2\frac{x + e_n}{[d_{Eucl}(x + e_n, 0)]^2} - e_n.$$

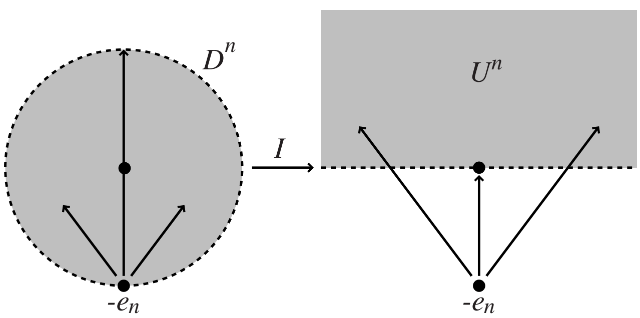

Figure 4.

Let $p_n\colon \mathbb{R}^n \to \mathbb{R}$ be projection onto the n^{th} co-ordinate. Let $x = (x_1, \ldots, x_n) \in \mathbb{R}^n - \{-e_n\}$. Then

$$I(x) \in U^n \Leftrightarrow p_n(I(x)) > 0$$
$$\Leftrightarrow p_n\left(2\frac{x + e_n}{[d_{Eucl}(x + e_n, 0)]^2} - e_n\right) > 0 \Leftrightarrow p_n\left(2\frac{x + e_n}{[d_{Eucl}(x + e_n, 0)]^2}\right) > 1$$
$$\Leftrightarrow \frac{2(x_n + 1)}{x_1^2 + \ldots + x_{n-1}^2 + (x_n + 1)^2} > 1$$

$$\Leftrightarrow x_1^2 + \ldots + x_n^2 < 1 \Leftrightarrow x \in D^n.$$

So I restricts to a diffeomorphism $D^n \to U^n$.

We now check that it is a Riemannian isometry. Note that $(DI)_x$ acts on $T_x D^n$ by scaling by a factor of $2/[d_{Eucl}(x+e_n, 0)]^2$, then reflecting in the direction of the line joining x to $-e_n$. So,

$$\begin{aligned}
||(DI)_x(v)||_{U^n} &= \frac{||(DI)_x(v)||_{Eucl}}{p_n(I(x))} \\
&= \left(\frac{2||v||_{Eucl}}{[d_{Eucl}(x+e_n,0)]^2}\right)\left(\frac{2(x_n+1)}{[d_{Eucl}(x+e_n,0)]^2} - 1\right)^{-1} \\
&= \frac{2||v||_{Eucl}}{2x_n + 2 - (x_1^2 + \ldots x_{n-1}^2 + (x_n+1)^2)} \\
&= \frac{2||v||_{Eucl}}{1 - [d_{Eucl}(x,0)]^2} \\
&= ||v||_{D^n}.
\end{aligned}$$

So, I is a Riemannian isometry. □

Note. The map I takes $S_\infty^{n-1} - \{-e_n\}$ to $\{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_n = 0\}$. Therefore, we view the 'sphere at infinity' for U^n to be the plane $\partial U^n = \{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_n = 0\}$, together with a single 'point at infinity', written ∞. In the case $n = 3$, this is the well-known observation that the Riemann sphere is just the complex plane with a single point added.

Note. The isometry $I \colon D^n \to U^n$ is a composition of Euclidean translations and scales, and the map $x \mapsto x/[d_{Eucl}(x, 0)]^2$, which is known as a *Euclidean inversion*.

Lemma 1.2. *I preserves the set $\{$Euclidean planes of dimension $k\} \cup \{$Euclidean spheres of dimension $k\}$.*

Proof. Since Euclidean scales and translations preserve this set, it suffices to show that Euclidean inversion does also. First consider the special case $k = n - 1$. Then, spheres and planes have the form

$$\{x \in \mathbb{R}^n : k_1 \langle x, x \rangle_{Eucl} + k_2 \langle x, a \rangle_{Eucl} + k_3 = 0\},$$

where k_1, k_2 and k_3 are real numbers, not all zero, and a is a vector in $\mathbb{R}^n$, and where these are chosen so that more than one x satisfies the equation. If we invert, this set is sent to the set

$$\begin{aligned}&\{x \in \mathbb{R}^n : k_1 \frac{\langle x, x\rangle_{Eucl}}{\langle x, x\rangle^2_{Eucl}} + k_2 \frac{\langle x, a\rangle_{Eucl}}{\langle x, x\rangle_{Eucl}} + k_3 = 0\}\\ =&\{x \in \mathbb{R}^n : k_1 + k_2\langle x, a\rangle_{Eucl} + k_3\langle x, x\rangle_{Eucl} = 0\},\end{aligned}$$

which, again is the equation of a sphere or plane in $\mathbb{R}^n$.

Now, consider the general case, where $1 \leq k \leq n-1$. Then a k-dimensional plane or sphere is the intersection of a collection of $(n-1)$-dimensional planes or spheres. This is mapped to the intersection of a collection of $(n-1)$-dimensional planes or spheres, which is an l-dimensional plane or sphere. Since I is a diffeomorphism, it preserves the dimension of submanifolds, and so $l = k$. □

The Klein model

Let $K^n = \{x \in \mathbb{R}^n : d_{Eucl}(x, 0) < 1\}$. Define

$$\begin{aligned}\phi\colon D^n &\to K^n\\ x &\mapsto x\left(\frac{2d_{Eucl}(x,0)}{[d_{Eucl}(x,0)]^2 + 1}\right).\end{aligned}$$

Assign K^n the metric that makes ϕ an isometry. This is the Klein model for $\mathbb{H}^n$. Unlike the other two models, angles in K^n do not agree with Euclidean angles. However, we will see that geodesics in K^n are Euclidean geodesics, after re-parametrisation.

2. Some isometries of $\mathbb{H}^n$.

Note that isometries $D^n \to D^n$ are in one-one correspondence with isometries $U^n \to U^n$, by conjugating by $I\colon D^n \to U^n$. Using this, we'll feel free to jump between the two different models we have for hyperbolic space.

Examples. 1. Any linear orthogonal map $h\colon \mathbb{R}^n \to \mathbb{R}^n$ fixing the origin restricts to an isometry $D^n \to D^n$. By considering $I \circ h \circ I^{-1}$ and noting that $I(0) = e_n$, we can find an isometry of U^n fixing e_n, and which realizes any orthogonal map $T_{e_n}U^n \to T_{e_n}U^n$.

2. Consider the map

$$\mathbb{R}^n \xrightarrow{h} \mathbb{R}^n$$

$$x \mapsto \lambda A x + b,$$

where $\lambda \in \mathbb{R}_{>0}$, A is an orthogonal map preserving the e_n-axis and $b \in \mathbb{R}^{n-1} \times \{0\}$. This restricts to a map $U^n \to U^n$ which is a hyperbolic isometry: if $x \in U^n$ and $v \in T_x U^n$, then

$$||(Dh)_x(v)||_{hyp} = ||\lambda v||_{hyp} = \frac{||\lambda v||_{Eucl}}{p_n(\lambda x)} = \frac{||v||_{Eucl}}{p_n(x)} = ||v||_{hyp}.$$

Theorem 2.1. *For any two points x and y in $\mathbb{H}^n$ and any orthogonal map $A \colon T_x\mathbb{H}^n \to T_y\mathbb{H}^n$, there is an isometry $h \colon \mathbb{H}^n \to \mathbb{H}^n$ such that $h(x) = y$ and $(Dh)_x = A$. Moreover, h is a composition of isometries as in Examples 1, 2 and 3.*

Proof. Consider x and y in U^n. By using Example 2, we may find isometries f and g such that $f(x) = g(y) = e_n$. Now, $(Dg)_y \circ A \circ (Df^{-1})_{e_n}$ is an orthogonal map $T_{e_n}U^n \to T_{e_n}U^n$ and so is realised by an isometry h fixing e_n, as in Example 1. Therefore, $g^{-1} \circ h \circ f$ is the required isometry. □

Definition. $\mathrm{Isom}(\mathbb{H}^n)$ is the group of isometries of $\mathbb{H}^n$. $\mathrm{Isom}^+(\mathbb{H}^n)$ is the subgroup of orientation-preserving isometries.

Corollary 2.2. *The isometries of Examples 1 and 2 generate* $\mathrm{Isom}(\mathbb{H}^n)$.

Proof. Suppose $h \in \mathrm{Isom}(\mathbb{H}^n)$. Pick $x \in \mathbb{H}^n$. By Theorem 2.1, there is an isometry $g \colon \mathbb{H}^n \to \mathbb{H}^n$ such that $g(x) = h(x)$ and $(Dg)_x = (Dh)_x$, with g a composition of isometries as in Examples 1, 2 and 3. By Theorem 1.5 of the Introduction of Riemannian manifolds, $h = g$. □

Corollary 2.3. *Any hyperbolic isometry $D^n \to D^n$ (respectively, $U^n \to U^n$)*

1. *extends to a homeomorphism $S^{n-1}_\infty \to S^{n-1}_\infty$ (respectively, $\partial U^n \cup \{\infty\} \to \partial U^n \cup \{\infty\}$),*
2. *preserves*

$$\{\text{Euclidean planes of dimension } k\} \cup \{\text{Euclidean spheres of dimension } k\},$$

3. *preserves the angles between S^{n-1}_∞ and arcs intersecting S^{n-1}_∞ (respectively, ∂U^n).*

Proof. These are all true for Examples 1 and 2. □

Let 0 be the origin in D^n.

Lemma 3.1. *For any point $x \in D^n - \{0\}$, the unit speed path α running along the Euclidean straight line L through 0 and x is a shortest path from 0 to x in the hyperbolic metric. Hence, it is a geodesic in D^n.*

Proof. Let $\alpha_1 \colon [0,T] \to D^n$ be another path from 0 to x in D^n. Our aim is to show that $\text{Length}_{hyp}(\alpha_1) \geq \text{Length}_{hyp}(\alpha)$. We may assume that $\alpha_1^{-1}(0) = 0$. Let α_2 be the path running along L such that

$$d_{Eucl}(\alpha_2(t), 0) = d_{Eucl}(\alpha_1(t), 0)$$

for all $t \in [0,T]$. Then,

$$||\alpha_2'(t)||_{Eucl} \leq ||\alpha_1'(t)||_{Eucl}.$$

Since $||\,||_{Eucl}$ and $||\,||_{hyp}$ differ by a factor which depends only on the Euclidean distance from 0, we have that

$$||\alpha_2'(t)||_{hyp} \leq ||\alpha_1'(t)||_{hyp}.$$

So,

$$\text{Length}_{hyp}(\alpha_2) = \int_0^T ||\alpha_2'(t)||_{hyp}\, dt \leq \int_0^T ||\alpha_1'(t)||_{hyp}\, dt = \text{Length}_{hyp}(\alpha_1),$$

But then $\alpha^{-1} \circ \alpha_2$ is a function $f \colon [0,T] \to [0, \text{Length}_{hyp}(\alpha)]$, such that $|f'(t)| = ||\alpha_2'(t)||_{hyp}$. Then

$$\text{Length}_{hyp}(\alpha_2) = \int_0^T ||\alpha_2'(t)||_{hyp}\, dt = \int_0^T |f'(t)|\, dt \geq$$

$$\int_0^T f'(t)\, dt = f(T) - f(0) = \text{Length}_{hyp}(\alpha).$$

Hence, α is a shortest path from 0 to x. □

Corollary 3.2. *The unit speed geodesic α in Lemma 3.1 is the unique geodesic between 0 and x (up to re-parametrisation).*

Proof. Suppose that α_1 is another geodesic between 0 and x. By Lemma 3.1, α_1 is a Euclidean straight line. Since it goes through x, $\alpha_1'(0)$ is a multiple of $\alpha'(0)$, and so α_1 is a re-parametrisation of α. □

Corollary 3.3. *Between any two distinct points in $\mathbb{H}^n$, there is a unique geodesic.*

Proof. Let x and y be distinct points in D^n. By Theorem 2.1, there is an isometry $h\colon D^n \to D^n$ which takes x to 0. This induces a bijection between geodesics through x and geodesics through 0. By Corollary 3.2, there is a unique geodesic between 0 and $h(y)$. So, there is a unique geodesic between x and y. □

Theorem 3.4. *Geodesics in D^n (respectively, U^n) are precisely the Euclidean straight lines and circles which hit S_∞^{n-1} (respectively, ∂U^n) at right angles.*

Proof. Let α be a path in D^n. Let x be a point on α. By Theorem 2.1, there is an isometry $h\colon D^n \to D^n$ such that $h(x) = 0$. Then

$$\begin{aligned}
&\alpha \text{ is a geodesic} \\
&\Leftrightarrow h(\alpha) \text{ is a geodesic} \\
&\Leftrightarrow h(\alpha) \text{ is a Euclidean straight line through } 0 \\
&\Leftrightarrow h(\alpha) \text{ is a Euclidean straight line or circle} \\
&\quad \text{through 0 which hits } S_\infty^{n-1} \text{ at right angles} \\
&\Leftrightarrow \alpha \text{ is a Euclidean straight line or circle hitting } S_\infty^{n-1} \text{ at right angles.}
\end{aligned}$$

Note that in the above, we used the fact that Euclidean circles through 0 do not hit S_∞^{n-1} at right angles. □

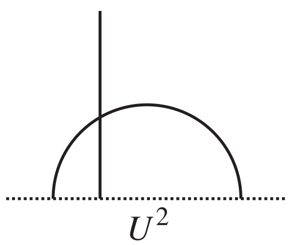

Figure 5.

Corollary 3.5. *If α is a geodesic in $\mathbb{H}^2$ and x is a point not on α, then there are infinitely many distinct geodesics through x which miss α.*

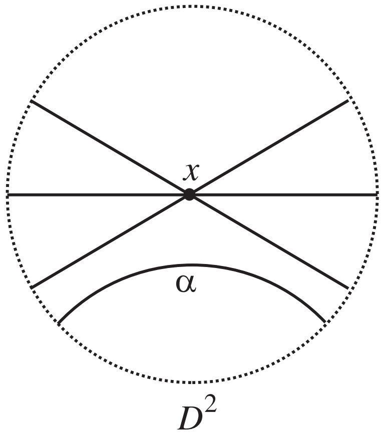

Figure 6.

Corollary 3.6. *Between any two distinct points x and y in S^{n-1}_∞, there is a unique geodesic.*

Proof. Work with U^n. After an isometry, we may assume that $x = \infty$. So $y \in \partial U^n$. By Theorem 3.4, geodesics through y are Euclidean straight lines and circles meeting y at right-angles. Therefore, the vertical straight line through y is the unique geodesic joining x to y. □

Corollary 3.7. *All geodesics in $\mathbb{H}^n$ are infinitely long in both directions.*

Proof. Let α be a geodesic in U^n. Perform an isometry of U^n so that it passes through e_n and runs to the point at ∞. Re-parametrise α so that the n^{th} coordinate of $\alpha(y)$ is y. Its Euclidean speed is then constant, but its hyperbolic speed is not. Its length between e_n and ∞ is

$$\int_1^\infty ||\alpha'(y)||_{hyp}\, dy = \int_1^\infty 1/y\, dy = [\ln(y)]_1^\infty = \infty. \qquad \square$$

Proposition 3.8. *Geodesics in the Klein model are Euclidean geodesics (up to re-parametrisation).*

Proof. We claim that the function $\phi\colon D^n \to K^n$ is described as follows. Embed D^n in D^{n+1}. At a point x of D^n, let γ be the geodesic perpendicular to D^n through x. Let $(x_1, \dots, x_n, \pm x_{n+1})$ be the endpoints of γ on S^n_∞. Then $\phi(x) = (x_1, \dots, x_n)$. This is a simple calculation in Euclidean geometry.

A codimension one hyperplane in D^n is described by an $(n-1)$-sphere intersecting S^{n-1}_∞ orthogonally. It is the intersection of D^n with an n-sphere S.

The vertical projection of $S \cap S^n_\infty$ to D^n is a Euclidean plane. Therefore ϕ maps codimension one hyperplanes in D^n to Euclidean planes. It therefore maps a hyperbolic geodesic, which is the intersection of hyperplanes, to a Euclidean geodesic. □

4. Classification of hyperbolic isometries

Recall the examples of hyperbolic isometries given in §2. The aim now is to show that every hyperbolic isometry is essentially one of these three types.

Definition. Let $h\colon \mathbb{H}^n \to \mathbb{H}^n$ be an isometry. Then h is

(i) *elliptic* if it fixes a point in $\mathbb{H}^n$;

(ii) *parabolic* if it has no fixed point in $\mathbb{H}^n$ and a unique fixed point in S^{n-1}_∞;

(iii) *loxodromic* if it has no fixed point in $\mathbb{H}^n$ and precisely two fixed points in S^{n-1}_∞.

Remarks. 1. Example 1 is elliptic. Example 2 is parabolic or elliptic if $\lambda = 1$, and is loxodromic if $\lambda \neq 1$.

2. If h and k are conjugate in $\mathrm{Isom}(\mathbb{H}^n)$, then h is elliptic (respectively, parabolic, loxodromic) if and only if k is elliptic (respectively, parabolic, loxodromic).

3. Some authors use the term 'hyperbolic' instead of loxodromic. This is confusing, since one can then talk about non-hyperbolic isometries of hyperbolic space.

Theorem 4.1. *Every isometry $h\colon \mathbb{H}^n \to \mathbb{H}^n$ is either elliptic, parabolic or loxodromic.*

We need to show two things: that every hyperbolic isometry has a fixed point somewhere either in hyperbolic space or in the sphere at infinity and that if it has no fixed in hyperbolic space, then it has at most two fixed points on the sphere at infinity.

Proposition 4.2. *Any isometry $U^n \to U^n$ fixing ∞ is of the form $x \mapsto \lambda A x + b$, where $\lambda \in \mathbb{R}_{>0}$, $b \in \mathbb{R}^{n-1} \times \{0\}$ and A is an orthogonal map fixing e_n.*

Proof. Let $h\colon U^n \to U^n$ be an isometry fixing ∞. It sends 0 to some point b. Let g be the translation $x \mapsto x+b$. Then $g^{-1}h$ fixes 0 and ∞. It therefore preserves the unique geodesic α between them. It acts as an isometry on α (which is isometric to $\mathbb{R}$). It maps some point x on α to λx. Let f be the map $x \mapsto \lambda x$. Then $f^{-1}g^{-1}h$ fixes x and acts on T_xU^n via some orthogonal map A fixing the e_n direction. This orthogonal map A is an isometry of U^n. By Theorem 1.5, $f^{-1}g^{-1}h$ equals A. Therefore, h is the map $x \mapsto \lambda Ax + b$. □

Corollary 4.3. *A non-elliptic isometry h which fixes at least two points on S_∞^{n-1} is conjugate to an isometry as in Example 2, with $\lambda \neq 1$. In particular, h is loxodromic.*

Proof. By conjugating the isometry, we may assume that the fixed points are at 0 and ∞. By Proposition 4.2, this isometry is of the form $x \mapsto \lambda Ax + b$, where b must be zero. If $\lambda = 1$, then it fixes all points on the e_n axis and hence h is elliptic. Therefore, $\lambda \neq 1$ and so the isometry is loxodromic. □

Proof of Theorem 4.1. Let h be a non-elliptic isometry. Corollary 4.3 implies it has most two fixed points on S_∞^{n-1}. We must show that it has at least one fixed point on S_∞^{n-1}. This is a special case of the Brouwer fixed point theorem which asserts that any continuous map from the closed unit ball in $\mathbb{R}^n$ to itself has a fixed point. Instead of quoting this, we'll prove it in this case.

Consider the displacement function

$$\mathbb{H}^n \xrightarrow{f} \mathbb{R}_{\geq 0}$$
$$x \mapsto d_{hyp}(h(x), x).$$

This is a continuous function. Either

1. the infimum of f is attained and is zero, or
2. the infimum of f is not attained, or
3. the infimum of f is attained and is non-zero.

Case 1. The point $x \in \mathbb{H}^n$ where the infimum is attained is a fixed point for h, and therefore h is elliptic.

Case 2. Then there is a sequence of points $x_1, x_2, \ldots$ in $\mathbb{H}^n$ such that $f(x_i) \to \inf(f)$. This sequence has a convergent subsequence in $\mathbb{H}^n \cup S_\infty^{n-1}$, since this is compact. Pass to this subsequence. The limit point x cannot lie in $\mathbb{H}^n$, since then $f(x) = \inf(f)$ and the infimum is attained. Therefore, $x \in S_\infty^{n-1}$. We will show that x is a fixed point for h.

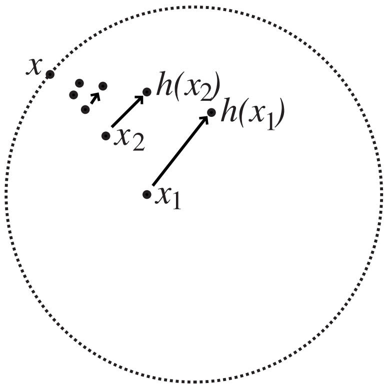

Figure 7.

Let I be the interval $[0, d_{hyp}(x_i, h(x_i))]$, and let $\alpha_i \colon I \to \mathbb{H}^n$ be the unit speed geodesic between x_i and $h(x_i)$. The sequence $d_{hyp}(x_i, h(x_i))$ is bounded above by some number M, say. Let 0 be the origin in D^n.

Since the points x_i are tending to a point on S_∞^{n-1}, the distance $d_{hyp}(x_i, 0) \to \infty$, by Corollary 3.7. By the triangle inequality,

$$d_{hyp}(\alpha_i(t), 0) \geq d_{hyp}(x_i, 0) - d_{hyp}(x_i, \alpha(t)) \geq d_{hyp}(x_i, 0) - M.$$

Therefore,

$$\inf_{t \in I} d_{hyp}(\alpha_i(t), 0) \to \infty$$

as $i \to \infty$. So,

$$\inf_{t \in I} d_{Eucl}(\alpha_i(t), 0) \to 1$$

as $i \to \infty$. Now,

$$1 = ||\alpha_i'(t)||_{hyp} = ||\alpha_i'(t)||_{Eucl} \left(\frac{2}{1 - [d_{Eucl}(\alpha_i(t), 0)]^2} \right).$$

So,

$$\sup_{t \in I} ||\alpha_i'(t)||_{Eucl} \to 0$$

as $i \to \infty$. Thus,

$$d_{Eucl}(x_i, h(x_i)) \leq \int_I ||\alpha_i'(t)||_{Eucl} \, dt \to 0$$

as $i \to \infty$. So, x is a fixed point of h.

Case 3. Let $x \in \mathbb{H}^n$ be a point where $\inf(f)$ is attained. We claim that the geodesic α through x and $h(x)$ is invariant under h. The endpoints of α on S_∞^{n-1} are therefore preserved. They are not permuted, since otherwise h would have a fixed point in α. Hence, h fixes these points on S_∞^{n-1}. This will prove the theorem.

Suppose that α is not invariant under h. The geodesics α and $h(\alpha)$ meet at $h(x)$. The angle between them is neither 0 or π (since that would imply that α was preserved by h).

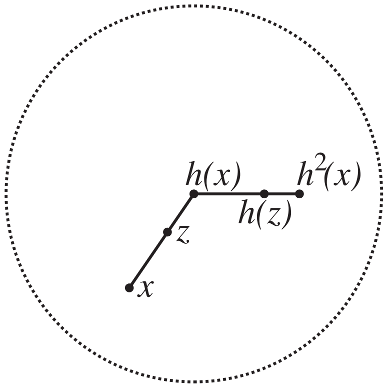

Figure 8.

Let z be any point on α between x and $h(x)$. Then

$$f(z) = d_{hyp}(z, h(z)) < d_{hyp}(z, h(x)) + d_{hyp}(h(x), h(z))$$

$$= d_{hyp}(z, h(x)) + d_{hyp}(x, z) = d_{hyp}(x, h(x)) = f(x).$$

Note that the inequality is strict because the angle between α and $h(\alpha)$ is not π. This contradicts the assumption that $\inf(f)$ is attained at x. □

We now know that every hyperbolic isometry is elliptic, parabolic or loxodromic. We also know from Corollary 4.3 that any loxodromic isometry is conjugate to one as in Example 2. What about elliptic and parabolic isometries?

Proposition 4.4. *An elliptic isometry is conjugate in* $\mathrm{Isom}(\mathbb{H}^n)$ *to an isometry as in Example 1.*

Proof. Let h be an elliptic isometry and let x be a fixed point for h in D^n. Pick an isometry k which takes x to 0. Then khk^{-1} fixes 0. It therefore acts on $T_0 D^n$

via an orthogonal map A. Let $A\colon D^n \to D^n$ be the isometry as in Example 1. By Theorem 1.5 of the Introduction to Riemannian manifolds, $khk^{-1} = A$. □

Proposition 4.5. *A parabolic isometry h is conjugate in* $\mathrm{Isom}(\mathbb{H}^n)$ *to an isometry as in Example 2, with $\lambda = 1$.*

Proof. By conjugating the isometry, we may assume that it fixes ∞ in the upper-half space model. By Proposition 4.2, it therefore acts as $x \mapsto \lambda Ax + b$ as in Example 2. We will show that $\lambda = 1$. Now, khk^{-1} is parabolic and so has no fixed point in ∂U^n. So,

$$x = \lambda Ax + b$$

has no solution. Therefore, $\det(\lambda A - I) = 0$, which means that λ^{-1} is a root of the characteristic polynomial for A. Hence, λ^{-1} is an eigenvalue of the orthogonal map A. So, $\lambda = 1$, since it is positive. □

Each isometry $\mathbb{H}^n \to \mathbb{H}^n$ extends to a homeomorphism $\mathbb{H}^n \cup S^{n-1}_\infty \to \mathbb{H}^n \cup S^{n-1}_\infty$. Therefore, this defines an *extension homomorphism*

$$\mathrm{Isom}(\mathbb{H}^n) \to \mathrm{Homeo}(S^{n-1}_\infty),$$

where $\mathrm{Homeo}(S^{n-1}_\infty)$ is the group of homeomorphisms of S^{n-1}_∞.

Proposition 4.6. *This homomorphism is injective.*

Proof. Suppose that an isometry h fixes S^{n-1}_∞. If h is elliptic, then, by Proposition 4.4, it is conjugate to an isometry as in Example 1. This must be the identity on $\mathbb{H}^n$, and therefore, h is the identity on $\mathbb{H}^n$.

Suppose now that h is non-elliptic. By Corollary 4.3, h fixes exactly two points on S^{n-1}_∞, which is a contradiction. □

What this means is that a hyperbolic isometry is determined by its action on the sphere at infinity.

5. PSL$(2, \mathbb{R})$ and PSL$(2, \mathbb{C})$

Definition. SL$(2, \mathbb{C})$ is the group of 2×2 matrices with entries in $\mathbb{C}$ and with determinant one. The group PSL$(2, \mathbb{C})$ is the quotient of SL$(2, \mathbb{C})$ by the normal subgroup $\{\mathrm{id}, -\mathrm{id}\}$. The groups SL$(2, \mathbb{R})$ and PSL$(2, \mathbb{R})$ are defined similarly.

There is a well-known relationship between PSL$(2, \mathbb{C})$ and Möbius maps. Associated with each element

$$\pm \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{PSL}(2, \mathbb{C}),$$

there is a Möbius map

$$\begin{aligned} \mathbb{C} \cup \{\infty\} &\to \mathbb{C} \cup \{\infty\} \\ z &\mapsto \frac{az+b}{cz+d}. \end{aligned}$$

This establishes an isomorphism between PSL$(2, \mathbb{C})$ and the group of Möbius maps (where the group operation in the latter is composition of maps).

Theorem 5.1. $\mathrm{Isom}^+(\mathbb{H}^3)$ *is isomorphic to* PSL$(2, \mathbb{C})$.

Identify ∂U^3 with $\mathbb{C}$, and S^2_∞ with $\mathbb{C} \cup \{\infty\}$. Consider the homomorphism

$$\mathrm{Isom}^+(\mathbb{H}^3) \to \mathrm{Homeo}(S^2_\infty) \cong \mathrm{Homeo}(\mathbb{C} \cup \infty),$$

which is injective by Proposition 4.6. Also, the map

$$\mathrm{PSL}(2, \mathbb{C}) \to \mathrm{Homeo}(\mathbb{C} \cup \{\infty\})$$

sending each matrix to its Möbius map is injective. Our aim is to show that the images of these two homomorphisms coincide.

Lemma 5.2. *The map* $\mathbb{C} \cup \{\infty\} \to \mathbb{C} \cup \{\infty\}$ *sending* $z \mapsto 1/z$ *is the extension of an orientation-preserving hyperbolic isometry.*

Proof. The isometry of D^3 will be the rotation ρ of angle π around the x-axis. This has the following effect on $\mathbb{C} \cup \{\infty\}$

$$\mathbb{C} \cup \{\infty\} \xrightarrow{I^{-1}} S^2_\infty \xrightarrow{\rho} S^2_\infty \xrightarrow{I} \mathbb{C} \cup \{\infty\}$$

sending a complex number z to a complex number z'. We must check that $z' = 1/z$. Clearly, $\arg(z') = -\arg(z)$.

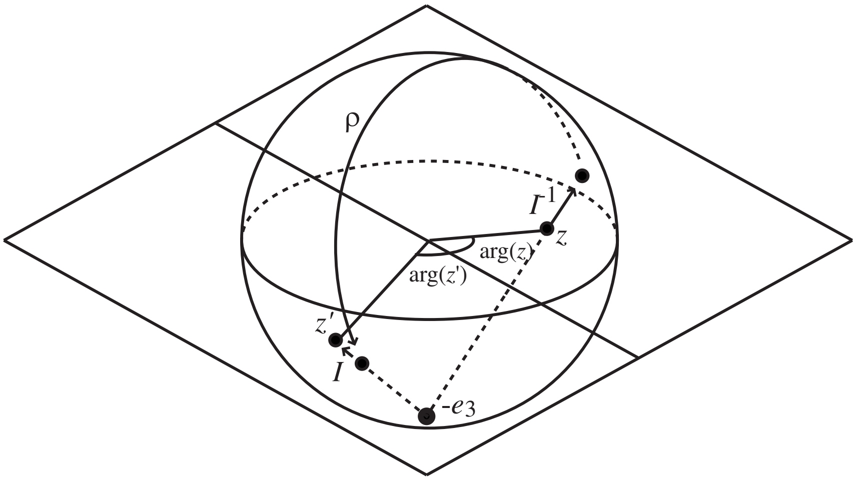

Figure 9.

We just have to check that $|z'| = 1/|z|$. Recall the definition of the map

$$\mathbb{R}^3 - \{-e_3\} \xrightarrow{I} \mathbb{R}^3 - \{-e_3\}$$
$$x \mapsto 2\frac{x+e_3}{[d_{Eucl}(x+e_3,0)]^2} - e_3.$$

Suppose that $x = (x_1, x_2, x_3) \in S^2_\infty$. So $x_1^2 + x_2^2 + x_3^2 = 1$. Also, $I(x) = (z_1, z_2, 0) = z \in \mathbb{C}$. Therefore

$$|z|^2 + 1 = [d_{Eucl}(z, -e_3))]^2 = 4[d_{Eucl}(x, -e_3)]^{-2}$$

$$= 4/[x_1^2 + x_2^2 + (x_3+1)^2] = 4/[2x_3+2] = 2/[x_3+1],$$

and so

$$|z|^2 = \frac{1-x_3}{1+x_3}.$$

Therefore, changing x_3 to $-x_3$ (which is the effect of ρ) changes $|z|^2$ to $|z|^{-2}$. Hence, $|z'| = |z|^{-1}$. □

Lemma 5.3. *Given any three distinct points x_1, x_2 and x_3 in $\mathbb{C} \cup \{\infty\}$, there is a Möbius map h such that $h(x_1) = 0$, $h(x_2) = 1$, $h(x_3) = \infty$.*

Proof. Consider the case where x_1, x_2 and x_3 are all in $\mathbb{C}$. Use the map

$$z \mapsto \left(\frac{x_2 - x_3}{x_2 - x_1}\right)\left(\frac{z - x_1}{z - x_3}\right).$$ □

Lemma 5.4. *Every element of* $\mathrm{Isom}^+(\mathbb{H}^3)$ *fixes some point on* S^2_∞.

Proof. This is true by definition if h is parabolic or loxodromic. If h is elliptic, then by Proposition 4.4, it is conjugate to an isometry as in Example 1. But any element of $SO(3)$ is a rotation which has at least two fixed points on S^2_∞. □

Proof of Theorem 5.1. We first show that every Möbius map is the extension of an orientation-preserving hyperbolic isometry. Any Möbius map can be expressed as a composition of the following maps

$$z \mapsto a_1 z$$

$$z \mapsto z + a_2$$

$$z \mapsto 1/z$$

The first and second of these are extensions of Example 2. Both of these are orientation-preserving. The third is an extension of an orientation-preserving elliptic isometry by Lemma 5.2.

We now show that every orientation-preserving hyperbolic isometry h extends to a Möbius map. By Lemma 5.4, h has a fixed point on S^2_∞. By Lemma 5.3, there is a Möbius map k sending this fixed point to ∞. So, khk^{-1} is an orientation-preserving isometry fixing ∞. So, it is of the form $z \mapsto \lambda Az + b$, as in Proposition 4.2. Since khk^{-1} is orientation-preserving, A is a rotation about e_n. So, khk^{-1} acts as $z \mapsto az + b$ ($a \in \mathbb{C} - \{0\}$, $b \in \mathbb{C}$) which is a Möbius map. □

Theorem 5.5. $\mathrm{Isom}^+(\mathbb{H}^2)$ *is isomorphic to* $\mathrm{PSL}(2, \mathbb{R})$.

Proof. Note that $\mathrm{PSL}(2, \mathbb{R})$ is the subgroup of $\mathrm{PSL}(2, \mathbb{C})$ which leaves $\mathbb{R} \cup \{\infty\} \subset \mathbb{C} \cup \{\infty\}$ invariant and preserves its orientation. Therefore, $\mathrm{PSL}(2, \mathbb{R})$ contains $\mathrm{Isom}^+(\mathbb{H}^2)$. To establish the opposite inclusion, we check that any orientation-preserving isometry of $\mathbb{H}^2$ extends to an orientation-preserving of $\mathbb{H}^3$. However, the orientation-preserving isometries of $\mathbb{H}^2$ are generated by the orientation-preserving isometries in Examples 1 and 2, and these extend to $\mathbb{H}^3$. □

6. Constructing hyperbolic manifolds

Recall that a hyperbolic manifold is a (G, X)-manifold, where X is $\mathbb{H}^n$ and $G = \mathrm{Isom}(\mathbb{H}^n)$. Since G is a group of Riemannian isometries, G acts rigidly. Hence, the pseudogroup generated by G is composed of diffeomorphisms between open subsets of $\mathbb{H}^n$ such that the restriction to any component is the restriction of an isometry of $\mathbb{H}^n$. This is so important that we include it as a proposition.

Proposition 6.1. *Let M be a hyperbolic manifold. Then M is a topological manifold with a cover by open sets U_i, together with open maps $\phi_i\colon U_i \to \mathbb{H}^n$ (known as charts) which are homeomorphisms onto their images, such that if $U_i \cap U_j \neq \emptyset$, then for each component X of $U_i \cap U_j$,*

$$\phi_j \circ \phi_i^{-1}\colon \phi_i(X) \to \phi_j(X)$$

is the restriction of a hyperbolic isometry.

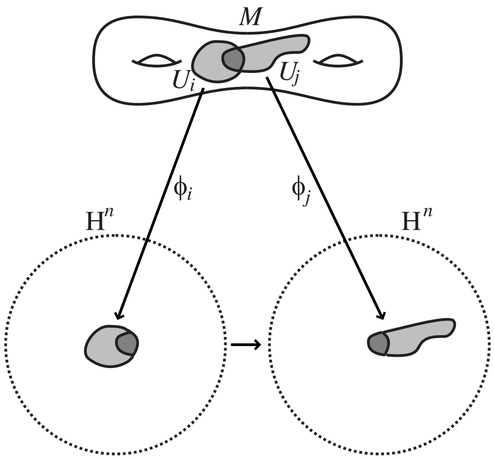

Figure 10.

Since G is a group of Riemannian isometries, a hyperbolic manifold inherits a Riemannian metric. The following characterises hyperbolic manifolds among all Riemannian manifolds.

Proposition 6.2. *Hyperbolic n-manifolds are precisely those Riemannian manifolds, each point of which has a neighbourhood isometric to an open subset of $\mathbb{H}^n$.*

Proof. In one direction this is straightforward. A hyperbolic manifold M has a chart into $\mathbb{H}^n$ which, by the way that the Riemannian metric on M is constructed, is an isometry to open subset of $\mathbb{H}^n$.

In the other direction, suppose that M is a Riemannian manifold, such that each point of M has a neighbourhood isometric to an open subset of $\mathbb{H}^n$. We take each such isometry to be a chart for M. We need to check that the transition maps lie in the pseudogroup generated by $\text{Isom}(\mathbb{H}^n)$. Suppose that $\phi_i: U_i \to \mathbb{H}^n$ and $\phi_j: U_j \to \mathbb{H}^n$ are charts that overlap. Let X be some component of $\phi_j(U_i \cap U_j)$. Then $\phi_i \circ \phi_j^{-1}|_X$ is a isometry between open subsets of $\mathbb{H}^n$. Let x be some point of X. Then, by Theorem 2.1, there is some isometry h of $\mathbb{H}^n$ such that $h(x) = \phi_i \circ \phi_j^{-1}(x)$ and $(Dh)_x = (D\phi_i \circ D\phi_j^{-1})_x$. By Theorem 1.5 of the Introduction to Riemannian Manifolds, $\phi_i \circ \phi_j^{-1}|_X$ is the restriction of h. □

Hyperbolic manifolds are obtained by gluing bits of hyperbolic space together via hyperbolic isometries. It is best to use 'nice' bits of hyperbolic space, for example, polyhedra.

Definition. A *k-dimensional hyperplane* in $\mathbb{H}^n$ is the image of $D^k \subset D^n$ after an isometry $D^n \to D^n$. A *half-space* is the closure in D^n of one component of the complement of a codimension one hyperplane.

Definition. A *polyhedron* in $\mathbb{H}^n$ is a compact subset of $\mathbb{H}^n$ that is the intersection of a finite collection of half-spaces. The *dimension* of a polyhedron is the smallest dimension of a hyperplane containing it. We will usually restrict attention to *non-degenerate* polyhedra in $\mathbb{H}^n$ which are those with dimension n. A *face* of a polyhedron P is the intersection $P \cap \pi$, where π is a codimension one hyperplane in $\mathbb{H}^n$ such that P is disjoint from one component of $\mathbb{H}^n - \pi$. Note that a face is a degenerate polyhedron. A *facet* is a face with codimension one. The *vertices* of P are the dimension zero faces. An *ideal polyhedron* is the intersection of a finite number of half-spaces in $\mathbb{H}^n$ whose closure in $\mathbb{H}^n \cap S^{n-1}_\infty$ intersects S^{n-1}_∞ in a finite number of points, and which has no vertices in $\mathbb{H}^n$.

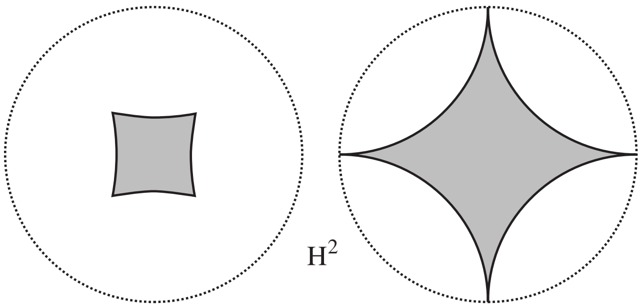

Figure 11.

Suppose now that M is obtained by gluing a collection of n-dimensional (possibly ideal) hyperbolic polyhedra $P_1, \dots, P_m$ by identifying their facets in pairs via isometries between the facets. Let P be the disjoint union $P_1 \cup \dots \cup P_m$, and let $q: P \to M$ be the quotient map. Note that $q|_{P - \partial P}$ is a homeomorphism and so $q(P - \partial P)$ inherits a hyperbolic structure. *Sometimes* M will be a hyperbolic manifold. But *sometimes* M may fail even to be a manifold. The following is a criterion which ensures that the hyperbolic structure on $q(P - \partial P)$ extends over all of M.

Theorem 6.3. *Suppose that each point $x \in M$ has a neighbourhood U_x and an open mapping $\phi_x: U_x \to B_{\epsilon(x)}(0) \subset D^n$ which is a homeomorphism onto its image, which sends x to 0 and which restricts to an isometry on each component of $U_x \cap q(P - \partial P)$. Then M inherits a hyperbolic structure.*

Proof. By reducing $\epsilon(x)$ if necessary, we can ensure that the closure of each component of $U_x \cap q(P - \partial P)$ contains x. Then ϕ_x will be the charts for M. These maps determine, for each $x_i \in q^{-1}(x)$, an isometry $h_{x_i}: B_{\epsilon(x)}(x_i) \to B_{\epsilon(x)}(0) \subset D^n$ such that $h_{x_i}|_{P - \partial P} = \phi_x \circ q$. We must check that if X is a component of $U_x \cap U_y$, then $\phi_y \phi_x^{-1}: \phi_x(X) \to \phi_y(X)$ is the restriction of a hyperbolic isometry. By assumption, this is true for each component of $\phi_x(X \cap q(P - \partial P))$. We must ensure that these isometries agree over all of $\phi_x(X)$. Any two points of $\phi_x(X \cap q(P - \partial P))$ are joined by a path in $\phi_x(X)$ which avoids (the image under $\phi_x \circ q$ of) the faces with dimension less than $n - 1$. Hence, we need only check that if z lies in $\phi_x(q(\partial P) \cap X)$ but not in a face of dimension less than $n - 1$, then then $\phi_y \phi_x^{-1}$ is an isometry in a neighbourhood of z. Let z_1 and z_2 be $q^{-1}\phi_x^{-1}(z)$. The component of $q^{-1}(U_x)$ (respectively, $q^{-1}(U_y)$) containing z_i contains a single point x_i of $q^{-1}(x)$

(respectively, y_i of $q^{-1}(y)$). Let F_i be the facet containing z_i and let $k\colon F_1 \to F_2$ be the identification isometry between the facets. Note that x_i and y_i also lie in F_i. There are two possible ways of extending k to an isometry of $\mathbb{H}^n$. Pick the one so that the following commutes:

$$\begin{array}{ccc} \mathbb{H}^n & \xrightarrow{k} & \mathbb{H}^n \\ & {\scriptstyle h_{x_1}} \searrow \quad \swarrow {\scriptstyle h_{x_2}} & \\ & D^n & \end{array}$$

Hence, running down the left-hand side of the following diagram is the same as running down the right-hand side (where the maps are defined):

$$\begin{array}{ccc} & D^n & \\ & {\scriptstyle h_{x_1}^{-1}} \swarrow \quad \searrow {\scriptstyle h_{x_2}^{-1}} & \\ \mathbb{H}^n & \xrightarrow{k} & \mathbb{H}^n \\ & {\scriptstyle h_{y_1}} \searrow \quad \swarrow {\scriptstyle h_{y_2}} & \\ & D^n & \end{array}$$

This ensures that $\phi_y \phi_x^{-1}$ is a well-defined isometry in a neighbourhood of z.

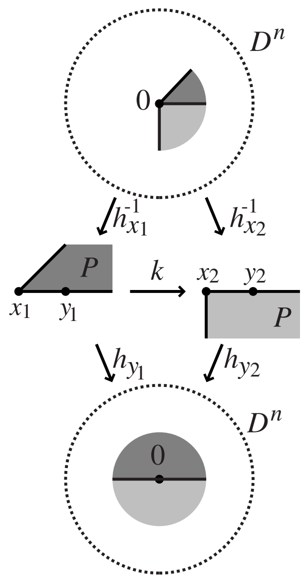

Figure 12.

There is one further thing to check: that M actually is a topological manifold. We must check that M is Hausdorff (which is clear) and that M has a countable basis of open sets. This is straightforward: we can refine $\{U_x : x \in M\}$ to a countable cover. Then, each U_x is homeomorphic to an open subset of $\mathbb{R}^n$, which has a countable basis of open sets. □

A similar theorem (with the same proof) also works for spherical and Euclidean structures.

This theorem reduces the problem of finding a hyperbolic structure to a problem one dimension lower. Define the *link* of a point x in a polyhedron P to be

$$\{v \in T_x\mathbb{H}^n : ||v|| = 1 \text{ and } \exp_x(\lambda v) \in P \text{ for some } \lambda > 0\}.$$

This is a polyhedron in the unit sphere in $T_x\mathbb{H}^n$. Isometries between facets of hyperbolic polyhedra induce isometries between facets of the links of identified points. The existence of a hyperbolic structure on the quotient space M is equivalent to each point of M having a link that is isometric to the $(n-1)$-sphere.

Here is a sample application of this theorem. Recall that every closed orientable 2-manifold is homeomorphic to one (and only one) of the surfaces F_0, F_1, ..., where $F_0 = S^2$ and F_k is obtained from F_{k-1} by removing the interior of an embedded 2-disc and then attaching a handle as follows:

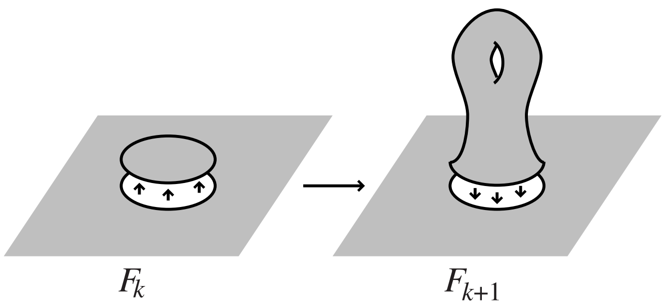

Figure 13.

Theorem 6.4. *F_k admits a hyperbolic structure for all $k \geq 2$.*

We shall see later that neither the 2-sphere nor the torus admits a hyperbolic structure.

Proof. For $k \geq 1$, F_k is obtained from a polygon with $4k$ sides by identifying their sides as follows:

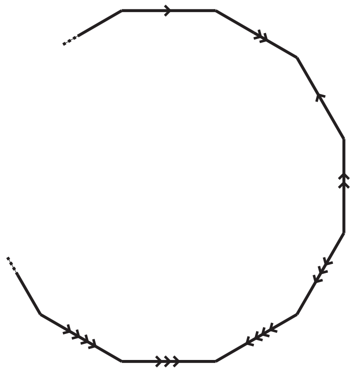

Figure 14.

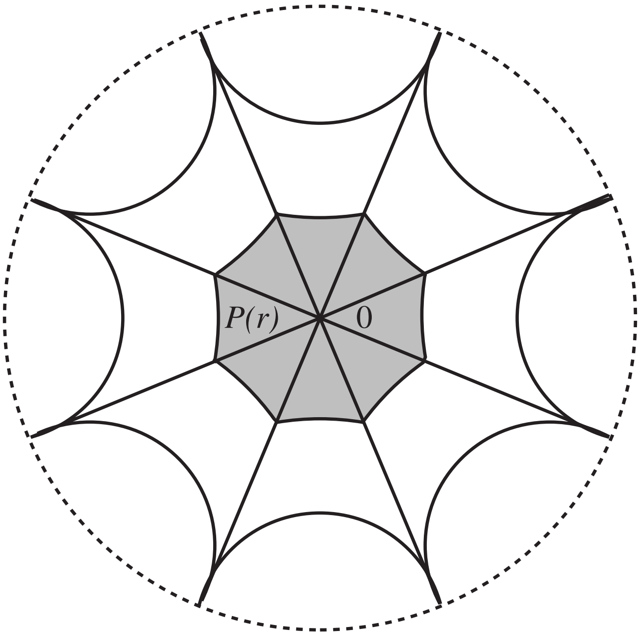

Figure 15.

Realise this as a polyhedron in $\mathbb{H}^2$, as follows. Draw $4k$ geodesics emanating from the origin 0 in D^2, with the angle between adjacent geodesics being $2\pi/4k$. Place a vertex on each geodesic, each a hyperbolic distance r from 0. Let $P(r)$ be the polyhedron having these $4k$ points as vertices. Let $\beta(r)$ be the interior angle at each vertex.

Claim. For some r, $\beta(r) = 2\pi/4k$.

Proof of theorem from claim. Glue the sides of $P(r)$ together to form F_k. Let

$q\colon P(r) \to F_k$ be the quotient map. We check that the requirements of Theorem 6.3 are satisfied, and hence that this determines a hyperbolic structure on F_k. If x lies in $q(P(r) - \partial P(r))$ then it automatically has a neighbourhood U_x as required.

If x lies in $q(\partial P(r))$ but is not in the image of a vertex of $P(r)$, then $q^{-1}(x)$ is two points. For each point y of $q^{-1}(x)$, $B_{\epsilon(x)}(y) \cap P(r)$ is isometric to half an $\epsilon(x)$-ball in D^n, providing $\epsilon(x)$ is sufficiently small. We may map these two half balls to a whole $\epsilon(x)$-ball in D^n as required.

There is a single point x lying in the image under q of the vertices. Since $\beta(r) = 2\pi/4k$, we may map $\epsilon(x)$-neighbourhoods of the $4k$ vertices homeomorphically to an $\epsilon(x)$-ball in D^n, in such a way that the homeomorphism restricts to an isometry on $q(P(r) - \partial P(r))$.

Proof of claim. Note that β is a continuous function of r. We examine the behaviour of $\beta(r)$ as $r \to 0$ and $r \to \infty$.

Perform a Euclidean scale h based at 0 which takes each vertex to a point on S^1_∞. This map preserves angles, but is *not* a hyperbolic isometry. As $r \to 0$, the sides of $h(P(r))$ approximate Euclidean straight lines. Hence, $\beta(r)$ tends to the interior angle of a regular Euclidean $4k$-gon, which is $\pi(1 - 1/2k)$.

As $r \to \infty$, the sides of $h(P(r))$ approximate hyperbolic geodesics. Hence, $\beta(r)$ tends to the angle between two geodesics emanating from the same point in S^1_∞. Since both geodesics are at right angles to S^1_∞, the angle between them is zero. Hence, $\beta(r) \to 0$ as $r \to \infty$.

Since $k > 1$, we have $0 < 2\pi/4k < \pi(1 - 1/2k)$. So, there is a value of r for which $\beta(r) = 2\pi/4k$. $\square$

7. The figure-eight knot complement

The following is a diagram of the figure-eight knot L. We will construct a hyperbolic structure on $S^3 - L$.

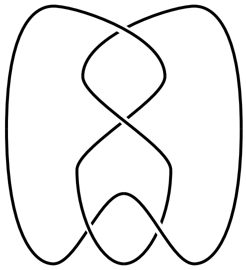

Figure 16.

Consider the following diagram of two regular tetrahedra in $\mathbb{R}^3$. There is a unique way to glue the faces in pairs (via Euclidean isometries) so that the edges (and their orientations) match. Let M be the space formed by this gluing.

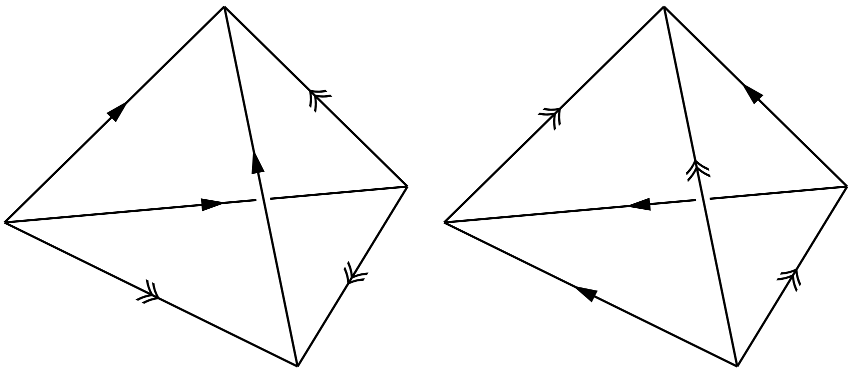

Figure 17.

This gives M a cell complex. However, M is not a manifold. The vertices of the tetrahedra are all identified to a single vertex v. A small neighbourhood of v is a cone on a torus.

However, every other point has a neighbourhood homeomorphic to an open ball in $\mathbb{R}^3$. So $M - v$ is a 3-manifold.

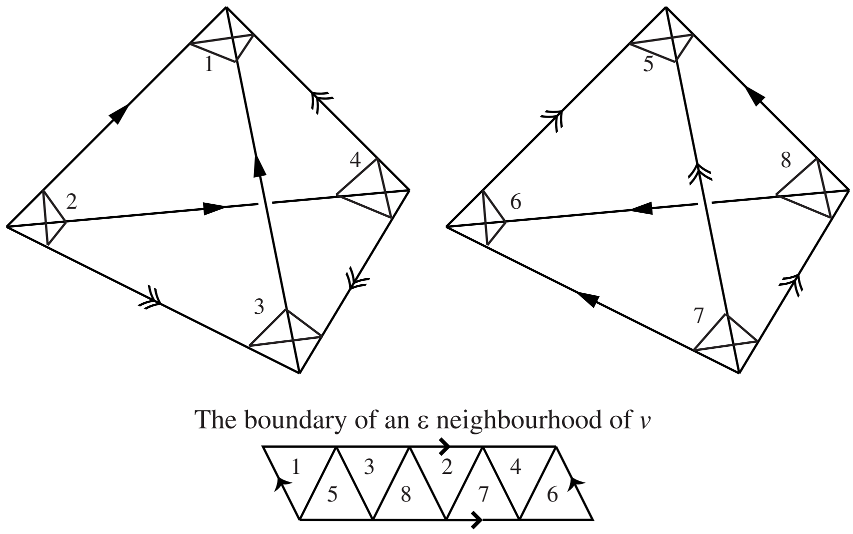

Figure 18.

Theorem 7.1. *$M - v$ is homeomorphic to $S^3 - L$, where L is the figure-eight knot.*

Proof. First consider the following cell complex K^1 embedded in S^3. The 1-cells are labelled 1 to 6 and are assigned an orientation.

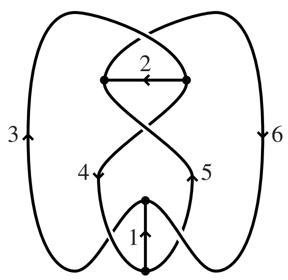

Figure 19.

Now attach to these four 2-cells, to form a cell complex K^2.

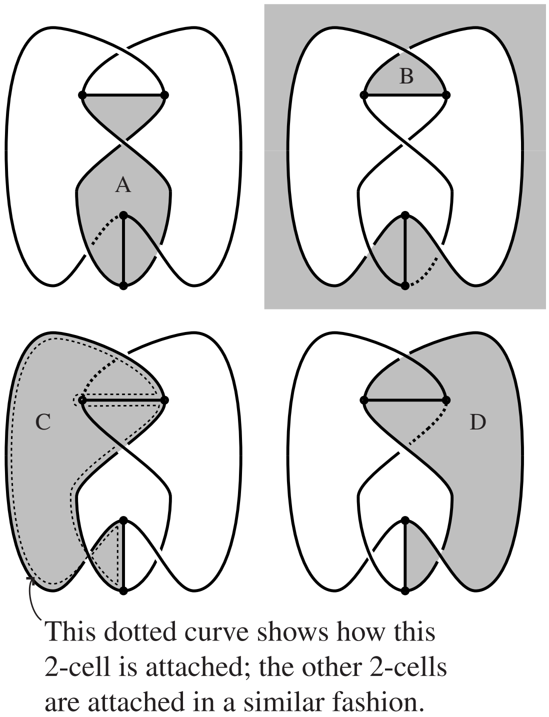

Figure 20.

Claim 1. $S^3 - K^2$ is homeomorphic to two 3-balls.

Before we prove this claim, the following will be useful.

Claim 2. There is a homeomorphism h between $S^3 - K^1$ and the complement of the following cell complex K_1^1.

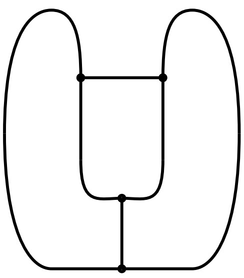

Figure 21.

The homeomorphism is supported in a neighbourhood of 1-cells 1 and 2. We focus on cell 1 (the other case being similar).

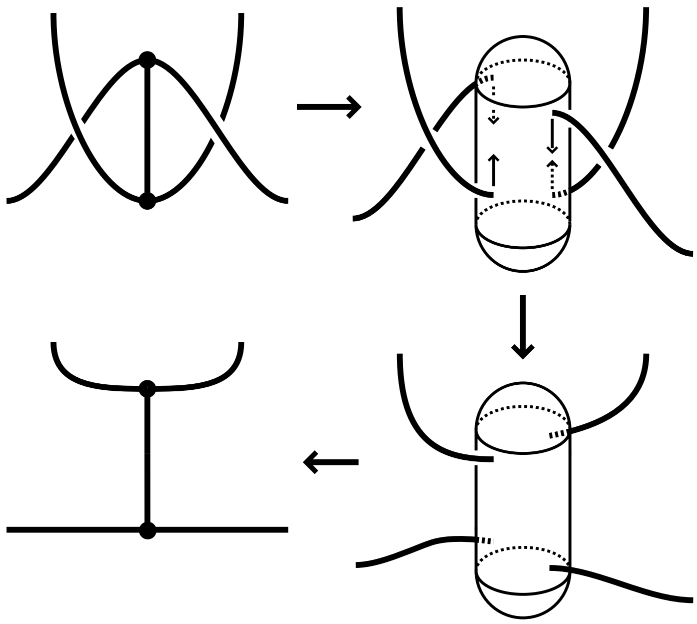

Figure 22.

In the first step above, we thicken cell 1 to a small closed 3-ball. In the second step, we slide the endpoints of the 1-cells attached to this ball. In the final step, we shrink the ball again to a single cell. This proves Claim 2.

The homeomorphism h takes the 2-cells of K^2 to the following 2-cells, which each lies in the plane of the diagram.

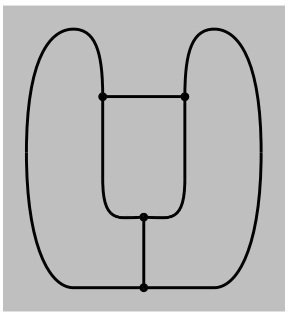

Figure 23.

Hence, $S^3 - K^2$ is homeomorphic to the complement of the above complex, which is two open 3-balls. This proves Claim 1.

Viewing these open 3-balls as the interiors of two 3-cells, we see that K^2

extends to a cell complex K^3 for S^3. The boundaries of the two 3-cells are attached onto K^2 according to the following diagram. Each 3-cell is a ball B glued onto K^2 via a map $\partial B \to K^2$. The following specifies this map for each 3-cell.

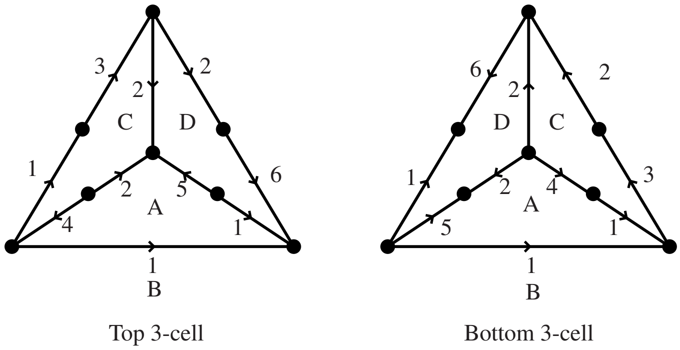

Figure 24.

The 0-cells and the 1-cells 3, 4, 5 and 6 combine to form the figure-eight knot L. If we collapse these cells to a point v, and then remove v, the result is the same as simply removing L from S^3. We therefore need to show that the cell complex we obtain by collapsing L to a point is M. It has a single 0-cell v, two 1-cells (coming from 1-cells 1 and 2), four 2-cells (coming from A, B, C and D), and two 3-cells. The attaching maps of the 2-cells and the 3-cells can be deduced from Figure 24, and are readily seen to give the required cell complex for M. Hence, $M - v$ is homeomorphic to $S^3 - L$. □

Definition. An *ideal n-simplex* is the ideal polyhedron determined by $n+1$ points on S^{n-1}_∞. An ideal 3-simplex is also known as an *ideal tetrahedron*.

Remark. If an ideal polyhedron is determined by some points V on S^{n-1}_∞, then we will often abuse terminology by calling V its *vertices*.

Definition. An ideal n-simplex is *regular* if, for any permutation of its vertices, there is a hyperbolic isometry which realises this permutation.

Construct a regular ideal tetrahedron as follows. Let V be the points in S^2_∞ which are the vertices of a regular Euclidean tetrahedron centred at the origin of D^3. Let Δ be the hyperbolic ideal tetrahedron determined by V. Then Δ is regular because any permutation of the points of V is realised by an orthogonal map of $\mathbb{R}^3$ which is hyperbolic isometry.

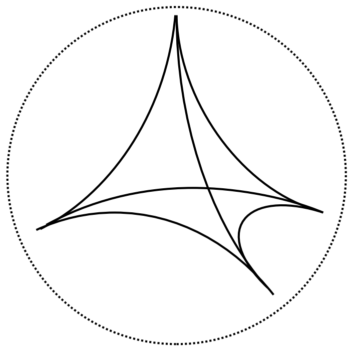

Figure 25.

Now glue two copies of Δ via isometries as specified by Figure 17. We will check that the conditions of Theorem 6.3 are satisfied and hence that this gives a hyperbolic structure on $S^3 - L$. As in Theorem 6.4, the only thing we have to check is that, for each point x lying in a 1-cell, the interior angles around x add up to 2π.

Note. The angle between two intersecting codimension one hyperplanes H_1 and H_2 in $\mathbb{H}^n$ is the same for all points of $H_1 \cap H_2$. This is because we may perform an isometry of D^n after which H_1 and H_2 both pass through 0. Then H_1 and H_2 become Euclidean hyperplanes, for which the assertion is clear.

Lemma 7.2. *Let F_1, F_2 and F_3 be three facets of an ideal tetrahedron in $\mathbb{H}^3$. Let β_{12}, β_{23} and β_{31} be the interior angles between these facets. Then $\beta_{12}+\beta_{23}+\beta_{31} = \pi$.*

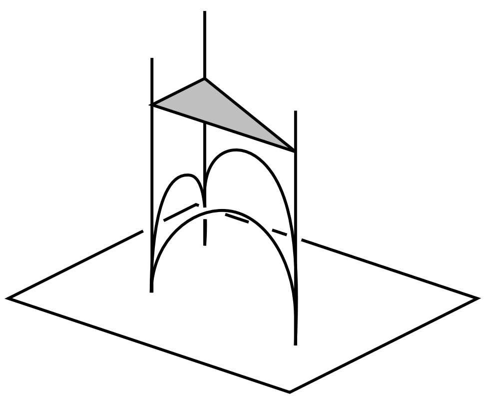

Figure 26.

Proof. Work with U^3. We may assume that the hyperplanes containing F_1, F_2 and F_3 are vertical Euclidean planes. Hence, β_{12}, β_{23} and β_{31} are the interior angles of a Euclidean triangle. □

Corollary 7.3. *If the ideal tetrahedron is regular, then* $\beta_{12} = \beta_{23} = \beta_{31} = \pi/3$.

For any point x of $M - v$ lying in a 1-cell, the 2-cells of M run past x six times. The six interior angles around x sum to 2π. Hence, we may construct the chart around x required by Theorem 6.3. Hence, $S^3 - L$ inherits a hyperbolic structure.

Remark. We have imposed a hyperbolic structure on the topological manifold $S^3 - L$. We have not shown that this is compatible with the smooth structure that $S^3 - L$ inherits from S^3 (although this is in fact true).

8. Gluing ideal tetrahedra

The construction of the hyperbolic structure on the figure-eight knot in the last section seems rather difficult to generalise. However, it is possible to construct many hyperbolic manifolds in this way, using ideal triangulations.

Definition. An *ideal triangulation* of a 3-manifold M is a way of constructing $M - \partial M$ from a collection of (topological) ideal tetrahedra by gluing their faces in pairs.

The following theorem (which we include without proof) demonstrates that there is no topological restriction to the existence of an ideal triangulation.

Theorem 8.1. *Any compact 3-manifold with non-empty boundary has an ideal triangulation.*

The aim of this section is to investigate when an ideal triangulation of a manifold can be used to impose a hyperbolic structure.

If the faces of an ideal tetrahedron in $\mathbb{H}^3$ are labelled 1 to 4, and β_{ij} is the angle between the faces i and j, then Lemma 7.2 gives the following equalities:

$$\beta_{23} + \beta_{34} + \beta_{42} = \pi$$
$$\beta_{34} + \beta_{41} + \beta_{13} = \pi$$
$$\beta_{41} + \beta_{12} + \beta_{24} = \pi$$
$$\beta_{12} + \beta_{23} + \beta_{31} = \pi$$

Adding the first two equations and subtracting the second two gives that $\beta_{12} = \beta_{34}$. Similarly, $\beta_{13} = \beta_{24}$ and $\beta_{14} = \beta_{23}$. Thus opposite edges have the same interior angle. An ideal tetrahedron in $\mathbb{H}^3$ therefore determines three interior angles α, β and γ adding up to π. All three angles appear at each vertex, and they cycle round the vertex the same way. Hence, if one considers such tetrahedra up to orientation-preserving isometry, the triple (α, β, γ) is well-defined up to cyclic permutation.

Lemma 8.2. *Ideal tetrahedra in* $\mathbb{H}^3$ *(up to orientation preserving isometry) are in one-one correspondence with triples of positive numbers adding to* π *(up to cyclic permutation).*

Proof. We have already seen how an ideal tetrahedron determines a triple. Any such triple can be realized by some ideal tetrahedron: find a Euclidean triangle in ∂U^3 with these interior angles and then consider the ideal tetrahedron with vertices being ∞ and the three corners of the triangle. Also, if two ideal triangles have the same interior angles, we can find an orientation-preserving isometry taking one to the other. Perform isometries taking one vertex of each tetrahedron to ∞. Then perform parabolic and loxodromic isometries that match up the remaining three vertices of each tetrahedron. □

There is an alternative way of describing these ideal tetrahedra. Given a triple (α, β, γ), consider a Euclidean triangle having these interior angles, and having vertices at 0, 1 and some point z in $\mathbb{C}$, where $\mathrm{Im}(z) > 0$. There is some ambiguity here since α, β or γ may be placed at the origin. Therefore, the following complex numbers all represent the same triple:

$$z, \ \frac{1}{1-z}, \ 1 - \frac{1}{z}.$$

Suppose now that M has an ideal triangulation, and that each ideal tetrahedron has been assigned interior angles α, β and γ. Each ideal tetrahedron then inherits a hyperbolic structure. The topological identification of facets of the tetrahedra is, according to the following lemma, realized by a unique isometry.

Lemma 8.3. *Let* Δ *and* Δ' *be ideal triangles in* $\mathbb{H}^2$ *with edges* (e_1, e_2, e_3) *and* (e'_1, e'_2, e'_3)*. Then there is a unique isometry of* $\mathbb{H}^2$ *taking* Δ *to* Δ' *and the edges* e_i *to edges* e'_i*.*

Proof. Send the vertex joining e_1 and e_2 to ∞ in the upper half space model. Also send the vertex joining e'_1 and e'_2 to ∞. Then there is a unique hyperbolic isometry taking the remaining pairs of vertices to each other. □

The hyperbolic structures on the tetrahedra then patch together to form a hyperbolic structure on the complement of the edges of M. When does this extend to a hyperbolic structure on M? The answer can be given in precise algebraic terms. At the i^{th} edge of M, let $w_{i1}, \dots, w_{ik}$ be the complex parameters of the tetrahedra around that edge. The choice of whether w_{ij} equals z, $1/(1-z)$ or $1 - 1/z$ is made so that the interior angle of the Euclidean triangle at the origin is the same as the interior angle at the edge i.

Theorem 8.4. *Let* M *be as above, with the extra condition that* ∂M *is a collection of tori. Then* $M - \partial M$ *inherits a hyperbolic structure if and only if for each edge* i

$$\prod_{j=1}^{k} w_{ij} = 1.$$

Proof. $M - \partial M$ is obtained from a collection of disjoint ideal tetrahedra in $\mathbb{H}^3$ by identifying pairs of faces. Let $e_1, \dots, e_k$ be the edges of these tetrahedra that are all identified to form the i^{th} edge e of M. We let $e_{k+1} = e_1$. The isometries between faces yield isometries $e_i \to e_{i+1}$. Hence we obtain an isometry

$$e_1 \to e_2 \to \dots \to e_1.$$

Case 1. $e_1 \to e_1$ is a non-zero translation along e_1.

Then each point on e has an infinite number of inverse-images in $e_1 \cup \dots \cup e_k$. So the quotient space is not a manifold.

Case 2. $e_1 \to e_1$ is a reflection.

This reflection has a fixed point. A neighbourhood of this point in the quotient space is a cone on $\mathbb{R}P^2$. Hence, the quotient space is not a manifold.

Case 3. $e_1 \to e_1$ is the identity.

In this case, the quotient space is indeed a manifold. We claim that this case holds if and only if

$$|\prod_{j=1}^{k} w_{ij}| = 1.$$

Place the tetrahedron containing e_1 in U^3 so that its vertices are ∞, 0, 1 and w_{i1}, and so that e_1 runs from 0 to ∞. Place the tetrahedron containing e_2 beside it, so that their faces are glued via the correct isometry. The vertices of this tetrahedra are ∞, 0, w_{i1} and $w_{i1}w_{i2}$. Continue this procedure for all k tetrahedra. In this way, we have ensured that all the gluing maps $e_i \to e_{i+1}$ are the identity for $1 \leq i < k$. The final gluing map $e_k \to e_1$ sends the points ∞, 0, $\prod_{j=1}^{k} w_{ij}$ to ∞, 0 and 1. It is therefore a loxodromy with invariant geodesic e_1. Thus $e_1 \to e_1$ is the identity if and only if $e_k \to e_1$ is the identity if and only if $|\prod_{j=1}^{k} w_{ij}| = 1$. This proves the claim.

Suppose that this condition holds. We can then apply Theorem 6.3. A chart into $\mathbb{H}^3$ exists at each point of each edge of $M - \partial M$ if and only if the angles around each edge sum to 2π. Thus, in summary, $M - \partial M$ inherits a hyperbolic structure if and only if

$$|\prod_{j=1}^{k} w_{ij}| = 1 \text{ for each } i \text{ and} \tag{1}$$

$$\sum_{j=1}^{k} \arg(w_{ij}) = 2\pi \text{ for each } i. \tag{2}$$

We need to show that this is equivalent to

$$\prod_{j=1}^{k} w_{ij} = 1 \text{ for each } i. \tag{3}$$

Clearly, (1) and (2) imply (3). Also, (3) implies (1). Also, (3) implies that

$$\sum_{j} \arg(w_{ij}) = 2\pi N(i),$$

for positive integers $N(i)$. We need to show that $N(i) = 1$ for each i. Summing the above inequalities over all i, and noting that each of the six interior angles in each tetrahedron appears exactly once gives that

$$2\pi T(M) = 2\pi \sum_{i=1}^{E(M)} N(i),$$

where $T(M)$ is the number of tetrahedra of M and $E(M)$ is the number of edges of M. Thus, $T(M) \geq E(M)$, with equality if and only if $N(i) = 1$ for each i. The ideal triangulation of M induces a triangulation of ∂M with $V(\partial M)$ vertices, $E(\partial M)$ edges and $F(\partial M)$ faces. Since ∂M is a collection of tori

$$\begin{aligned} 0 = \chi(\partial M) &= V(\partial M) - E(\partial M) + F(\partial M) \\ &= V(\partial M) - 3F(\partial M)/2 + F(\partial M) \\ &= V(\partial M) - F(\partial M)/2 \\ &= 2E(M) - 2T(M). \end{aligned}$$

Hence, $E(M) = T(M)$. So, $N(i) = 1$ for each i. Therefore, (3) implies (2). □

We now apply Theorem 8.4 to the case where M is the exterior of the figure-eight knot. We assign complex numbers z_1 and z_2 to the two tetrahedra. There are two edges of M, giving the following equations:

$$1 = z_2 z_1 \left(1 - \frac{1}{z_2}\right) z_1 z_2 \left(1 - \frac{1}{z_1}\right)$$
$$1 = \left(\frac{1}{1-z_2}\right)\left(1 - \frac{1}{z_1}\right)\left(\frac{1}{1-z_2}\right)\left(\frac{1}{1-z_1}\right)\left(1 - \frac{1}{z_2}\right)\left(\frac{1}{1-z_1}\right).$$

These equations are equivalent, since their product is $1 = 1$. To see this, note that

$$z\left(\frac{1}{1-z}\right)\left(1 - \frac{1}{z}\right) = -1.$$

These equations can be written more neatly as

$$z_1(z_1 - 1)z_2(z_2 - 1) = 1.$$

The hyperbolic structure imposed on M in §7 was the case where $z_1 = z_2 = e^{i\pi/3}$. However, there is a one-complex-dimensional parametrisation of hyperbolic structures that arise by perturbing z_1 and z_2 from this value in such a way that the equation $z_1(z_1 - 1)z_2(z_2 - 1) = 1$ remains satisfied.

9. Completeness

We will prove the following theorem over the next few sections.

Theorem 9.1. *If M is a simply-connected complete hyperbolic n-manifold, then M is isometric to $\mathbb{H}^n$.*

Note that the any cover $\tilde{M}$ of a Riemannian manifold M inherits a Riemannian metric from M. Any covering transformations of $\tilde{M}$ are isometries. Note also that $\tilde{M}$ is complete if M is complete.

Corollary 9.2. *The universal cover of a complete hyperbolic n-manifold is isometric to $\mathbb{H}^n$.*

Theorem 9.1 is a special case of the following theorem, which we will prove over the next two sections.

Theorem 9.3. *Let G be a group of isometries of a simply-connected Riemannian manifold X. Let M be a complete simply-connected (G, X)-manifold. Then M is (G, X)-isomorphic (and hence isometric) to X.*

The hypothesis that M is complete comes in via the following well known result from Riemannian geometry.

Theorem 9.4. [Hopf-Rinow] *Let M be a connected Riemannian manifold. Then the following are equivalent:*

1. *M is complete as a metric space;*
2. *any geodesic $I \to M$ can be extended to a geodesic $\mathbb{R} \to M$;*
3. *for any $m \in M$, $\exp_m$ is defined on all of $T_m M$;*
4. *for some $m \in M$, $\exp_m$ is defined on all of $T_m M$;*
5. *any closed bounded subset of M is compact.*

The proof requires machinery from differential geometry that we have not developed. However, all we need is that (1) $\Rightarrow$ (2) when M is hyperbolic. This we now prove.

Proof. Suppose that M is a complete hyperbolic manifold. For each non-zero vector $v \in T_x M$, let $\alpha\colon I \to M$ be the geodesic with $\alpha(0) = x$ and $\alpha'(0) = v$, and where $I \subset \mathbb{R}$ is the maximal domain of definition of α. It is a consequence of Proposition 1.3 in the Introduction to Riemannian Manifolds that I is an open neighbourhood of 0. We will show that the completeness of M implies that I is closed and hence the whole of $\mathbb{R}$. Let t_i be a sequence of points in I, converging to some point $t_\infty \in \mathbb{R}$. This is a Cauchy sequence in $\mathbb{R}$. The fact that $||\alpha'(t)||$ is constant implies that $\alpha(t_i)$ is Cauchy in M. Hence, it converges to a point $y \in M$. Pick a chart $\phi: U \to D^n$ around y, where $\phi(y) = 0 \in D^n$. Then $\phi \circ \alpha$ is a Euclidean straight line approaching 0. Hence, it can be smoothly extended. □

The following is a useful way of checking completeness.

Proposition 9.5. *Let M be a metric space. Suppose that there is some family of compact subsets S_t of M (for $t \in \mathbb{R}_{>0}$) which cover M, such that S_{t+a} contains all points within distance a of S_t. Then M is complete.*

Proof. Any Cauchy sequence in M must be contained in S_t for some sufficiently large t. Hence, it converges, since S_t is compact. □

Corollary 9.6. $\mathbb{H}^n$ *is complete.*

Proof. Let $S_t = B_t(0) \subset D^n$. Then $B_t(0)$ is a closed Euclidean ball. (Here, we are implicitly applying Corollary 3.7 to geodesics through 0.) So, S_t is compact. □

Example. In the case of the hyperbolic structure on the figure-eight knot complement defined in §7, let $S_t \cap \Delta$ for each of the two ideal tetrahedra Δ be $B_t(0) \cap \Delta$, where 0 is the origin in D^3. These S_t satisfy the condition of Proposition 9.5. Hence, the hyperbolic structure is complete. Applying Corollary 9.2 to this case, we obtain the following purely topological corollary.

Corollary 9.7. *The universal cover of the figure-eight knot complement is homeomorphic to* $\mathbb{R}^3$.

According to Thurston's theorem, the complements of 'most' knots in S^3 admit a complete finite volume hyperbolic structure. Therefore, their universal covers are all homeomorphic to $\mathbb{R}^3$.

Example. Here is an example of an incomplete hyperbolic structure. Let

$$B = \{(x_1, x_2) \in U^2 : 1 \leq x_1 \leq 2\}.$$

Glue the two sides of B via the isometry $z \mapsto 2z$. The resulting space M is homeomorphic to $S^1 \times \mathbb{R}$, and inherits a hyperbolic structure. This is incomplete: here is a Cauchy sequence which does not converge. Let $z_i = (1, 2^i) \in U^2$. This is identified with the point $(2, 2^{i+1})$. So,

$$d(z_i, z_{i+1}) \leq d_{hyp}((2, 2^{i+1}), (1, 2^{i+1})) < 1/2^{i+1}.$$

So, this sequence is Cauchy in M. However, it does not converge to a point in M, since the x_2 co-ordinates tend to ∞.

Theorem 9.3 is proved by defining a local (G, X)-isomorphism (and hence a local isometry) from M to X. The existence of this local isometry together with the following proposition will prove Theorem 9.3.

Proposition 9.8. *Let $h \colon M \to N$ be a local isometry between Riemannian manifolds, where M is complete. Then h is a Riemannian covering map.*

Proof. Let x be any point in N, and let y_i be the points of $h^{-1}(x)$. By Proposition 1.3 of the Introduction to Riemannian Manifolds, there is some r such that $\exp_x$ maps $B_r(0)$ diffeomorphically onto $B_r(x)$. Let $U_i = \exp_{y_i}(B_r(0))$. This is well-defined since $\exp_{y_i}$ is defined on all of $T_{y_i}M$ by the Hopf-Rinow theorem. Recall from Proposition 1.4 of the Introduction to Riemannian Manifolds that the following diagram commutes:

$$\begin{array}{ccc} T_{y_i}M & \xrightarrow{(Dh)_{y_i}} & T_xN \\ \big\downarrow \exp_{y_i} & & \big\downarrow \exp_x \\ M & \xrightarrow{h} & N \end{array}$$

Claim. $h|_{U_i}$ is a diffeomorphism onto its image.

The top and right arrows in the above diagram are diffeomorphisms when restricted to $B_r(0)$. Hence the bottom arrow must be a diffeomorphism when restricted to U_i.

Claim. $U_i \cap U_j = \emptyset$ if $i \neq j$.

If U_i and U_j overlap at a point m, there are vectors v_i and v_j in $T_{y_i}M$ and $T_{y_j}M$ with length at most r such that $m = \exp_{y_i}(v_i) = \exp_{y_j}(v_j)$. Using the above commutative diagram and the fact that $\exp_x$ is injective on $B_r(0)$, we deduce that

$(Dh)_{y_i}(v_i) = (Dh)_{y_j}(v_j)$. Hence, for all $t > 0$, $h \circ \exp_{y_i}(tv_i) = h \circ \exp_{y_j}(tv_j)$. By considering t near 1 and using the fact that h is injective near m, the geodesics $\exp_{y_i}(tv_i)$ and $\exp_{y_j}(tv_j)$ have the same derivative at $t = 1$. Hence they agree for all t. In particular, $y_i = \exp_{y_i}(0) = \exp_{y_j}(0) = y_j$. This proves the claim.

Claim. $\bigcup_i U_i = h^{-1}(B_r(x))$

Clearly, $\bigcup_i U_i \subset h^{-1}(B_r(x))$. To establish the opposite inclusion, consider a point $m \in M$ that is sent to a point in $B_r(x)$. Some geodesic of length $l < r$ runs from $h(m)$ to x, with velocity vector v at $h(m)$. There is a geodesic emanating from m with derivative $(Dh)_m^{-1}(v)$. Since M is complete, the Hopf-Rinow theorem gives that this geodesic is still defined after time l. By that stage it has reached some y_i. Hence, m lies in U_i.

This sequence of claims establishes that h is a covering map. □

10. The developing map

The goal of this section is prove Theorem 9.3. In fact, we will prove the following stronger result.

Theorem 10.1. *Let G be a group acting rigidly on a manifold X. Let M be a simply-connected (G, X)-manifold. Then any connected chart $\phi: U_0 \to X$ can be extended to a local (G, X)-isomorphism $D: M \to X$, known as a developing map.*

Note that in the above, we did not assume that X was a Riemannian manifold, complete or incomplete.

Proof. Start with a chart $\phi_0: U_0 \to X$, where U_0 is connected. Pick a basepoint $x_0 \in U_0$. We wish to define $D(x)$ for all points $x \in M$. There is a path $\alpha: [0, 1] \to M$ from x_0 to x. We will now define a corresponding path $\beta: [0, 1] \to X$, such that $\beta = \phi_0 \circ \alpha$ wherever the maps are defined. In particular, $\beta(0) = \phi_0(x_0)$. We will define $D(x)$ as $\beta(1)$.

For each point $t \in [0, 1]$, pick a chart around $\alpha(t)$. The inverse image of these charts forms an open cover of $[0, 1]$. By replacing each set with its connected components, we obtain a new open cover of $[0, 1]$. Since $[0, 1]$ is compact, there is a finite subcover $\{(t_0^-, t_0^+) \dots, (t_m^-, t_m^+)\}$ of $[0, 1]$. The interval (t_i^-, t_i^+) is open, half-open or closed in $\mathbb{R}$ as appropriate. Each (t_i^-, t_i^+) is a connected component of $\alpha^{-1}(U_i)$ for some chart $\phi_i \colon U_i \to X$. We may assume that (t_0^-, t_0^+) is the component of $\alpha^{-1}(U_0)$ containing 0. We may also assume that, if $i < j$, then $t_i^- \leq t_j^-$ and $t_i^+ \leq t_j^+$. This implies that $t_i^+ > t_{i+1}^-$ for all i.

The interval $(t_i^-, t_i^+) \cap (t_{i+1}^-, t_{i+1}^+)$ maps to a path in M, lying entirely in a connected component V_i of $U_i \cap U_{i+1}$. This has an associated transition map $\phi_{i+1} \circ \phi_i^{-1}$, which is the restriction of an element $h_i \colon X \to X$ of G. For $t \in [0,1]$, we pick (t_i^-, t_i^+) containing t and define $\beta(t)$ as

$$\beta(t) = h_0^{-1} \circ \ldots \circ h_{i-1}^{-1} \circ \phi_i \circ \alpha(t).$$

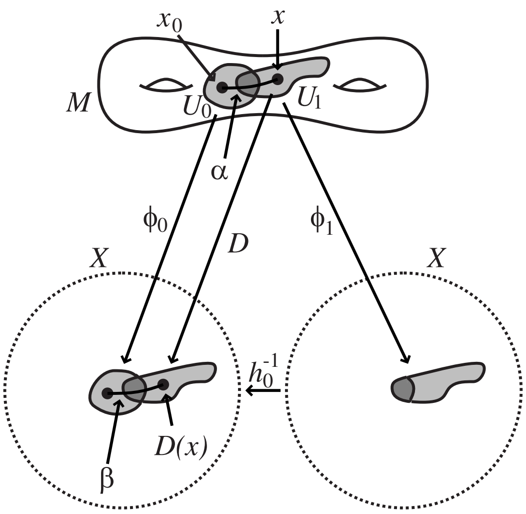

Figure 28.

We now show that this is independent of the choice of (t_i^-, t_i^+) containing t. If $t \in (t_k^-, t_k^+)$, with $k > i$, then $t \in (t_j^-, t_j^+)$ for all $i \leq j \leq k$. Then $\alpha(t)$ lies in $V_i, V_{i+1}, \ldots, V_{k-1}$. On these sets, the following maps are equal: $\phi_i = h_i^{-1} \circ \phi_{i+1}$, ..., $\phi_{k-1} = h_{k-1}^{-1} \circ \phi_k$. Hence,

$$h_0^{-1} \circ \ldots \circ h_{i-1}^{-1} \circ \phi_i \circ \alpha(t) = h_0^{-1} \circ \ldots \circ h_{k-1}^{-1} \circ \phi_k \circ \alpha(t).$$

We now show that β is independent of the choice of open cover of $[0,1]$. Suppose that $\{(\hat{t}_0^-, \hat{t}_0^+), \ldots, (\hat{t}_{\hat{m}}^-, \hat{t}_{\hat{m}}^+)\}$ is another cover of $[0,1]$, with each $(\hat{t}_i^-, \hat{t}_i^+)$ being a connected component of $\alpha^{-1}(\hat{U}_i)$ for some chart $\hat{\phi}_i \colon \hat{U}_i \to X$. Let $\hat{h}_i$, $\hat{V}_i$ and $\hat{\beta}$ be as above. We claim that the maps

$$\phi_0$$

$$h_0^{-1} \circ \phi_1$$

$$h_0^{-1} \circ h_1^{-1} \circ \phi_2$$

$$\dots$$

$$h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m$$

and

$$\phi_0$$

$$\hat{h}_0^{-1} \circ \hat{\phi}_1$$

$$\hat{h}_0^{-1} \circ \hat{h}_1^{-1} \circ \hat{\phi}_2$$

$$\dots$$

$$\hat{h}_0^{-1} \circ \ldots \circ \hat{h}_{m-1}^{-1} \circ \hat{\phi}_m$$

are equal on sets on which they are defined. If not, there is an infimal value of t ($t_{\inf}$, say) such that $t \in (t_i^-, t_i^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$ and

$$h_0^{-1} \circ \ldots h_{i-1}^{-1} \circ \phi_i \neq \hat{h}_0^{-1} \circ \ldots \hat{h}_{k-1}^{-1} \circ \hat{\phi}_k.$$

Then there is a $t < t_{\inf}$ such that $t \in (t_{i-1}^-, t_{i-1}^+) \cap (t_i^-, t_i^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$ or $(t_i^-, t_i^+) \cap (\hat{t}_{k-1}^-, \hat{t}_{k-1}^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$. Suppose the former. Then,

$$h_0^{-1} \circ \ldots h_{i-2}^{-1} \circ \phi_{i-1} = \hat{h}_0^{-1} \circ \ldots \hat{h}_{k-1}^{-1} \circ \hat{\phi}_k.$$

But

$$h_0^{-1} \circ \ldots h_{i-2}^{-1} \circ \phi_{i-1} = h_0^{-1} \circ \ldots h_{i-1}^{-1} \circ \phi_i.$$

The crucial fact we are using here is that if two elements of G agree on an open set, then they are equal, since G acts rigidly.

Now define $D(x)$ as $\beta(1)$. We need to show that this is independent of the choice of α. If $\hat{\alpha}$ is another path from x_0 to x, then there is a homotopy $H\colon [0,1] \times [0,1] \to M$ between α and $\hat{\alpha}$, keeping their endpoints fixed, since M is simply-connected. Pick a cover $\mathcal{C}$ of charts for M, one being U_0, the remainder being connected open sets disjoint from x_0. Subdivide $[0,1]$ into intervals,

$[0/N, 1/N], [1/N, 2/N]$, etc. Since $[0,1] \times [0,1]$ is compact, we may pick N large enough so that each square $[y/N, (y+1)/N] \times [z/N, (z+1)/N]$ lies within $H^{-1}(U)$ for some U of $\mathcal{C}$. So, $H([y/N, (y+1)/N] \times [z/N, (z+1)/N])$ lies within a compact subset K of U. Hence, there is a collection of paths $\alpha = \alpha_0, \alpha_1, \ldots, \alpha_k = \hat{\alpha}$, such that α_j and α_{j+1} differ only by a homotopy which alters the path within a compact subset K of U, keeping their endpoints fixed. By removing K from all the charts of $\mathcal{C}$ other than U, we may assume that that U is one of the charts U_i in the definition of β_j and β_{j+1}, and that the sets $U_0, \ldots, U_m, V_0, \ldots, V_{m-1}$ are the same for both α_j and α_{j+1}. So, β_j and β_{j+1} differ by a homotopy keeping their endpoints fixed. This does not alter $D(x)$. Thus, D is a well-defined map.

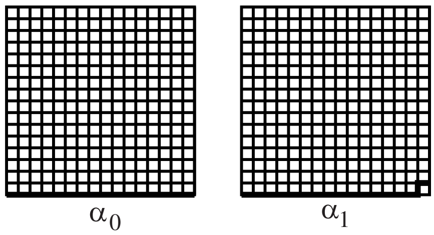

Figure 29.

We now show that D is a local (G, X)-isomorphism in a neighbourhood of each point $x \in M$. By definition, $D|_{U_0} = \phi_0$. So, we may assume that $x \neq x_0$. Pick a connected chart U_m around x not containing x_0, and let U be a connected open neighbourhood of x whose closure lies in U_m. Cover M by charts, one of which is U_m, the remainder being disjoint from U. Pick any path α from x_0 to x. Then, using the above cover,

$$D(x) = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m \circ \alpha(1) = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m(x),$$

where $h_i \colon X \to X$ are the relevant transition maps and $\phi_m \colon U_m \to \mathbb{H}^n$ is the chart. Define $D|_U$ using extensions of α within U. Then, $D|_U = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m$. Hence, $D|_U$ is a (G, X)-isomorphism onto its image. □

Example. Let $B = \{(x_1, x_2) \in U^2 : 1 \leq x_1 \leq 2\}$ and let M be the incomplete hyperbolic manifold obtained by gluing the two sides of B via the isometry $z \mapsto 2z$.

Let $p: \tilde{M} \to M$ be the universal cover. The sets

$$A_0 = \{(x_1, x_2) \in U^2 : 1 < x_1 < 2\}$$
$$A_1 = \{(x_1, x_2) \in U^2 : 1 \leq x_1 < 3/2\}$$
$$\cup \{(x_1, x_2) \in U^2 : 3/2 < x_1 \leq 2\}$$

form charts for M. These lift to charts for $\tilde{M}$. Let $\phi_0: U_0 \to A_0$ be the initial chart. Then, the associated developing map has image $\{(x_1, x_2) \in U^2 : 0 < x_1\}$.

The following lemma implies that, given the initial chart $\phi_0: U_0 \to X$, the developing map is unique.

Lemma 10.2. *Let G be a group acting rigidly on a space X. Let f_1 and f_2 be local (G, X)-isomorphisms between (G, X)-manifolds M and N, where M is connected. If f_1 and f_2 agree on some open set, then $f_1 = f_2$.*

Proof. Consider the set

$$V = \{x \in M : f_1 \text{ and } f_2 \text{ agree in some open neighbourhood of } x\}.$$

Clearly, V is open. We will now show that it is closed and so the whole of M. Consider a sequence of points $x_i \in V$ tending to x_∞. There are connected charts $\phi_M: U_M \to X$ and $\phi_N: U_N \to X$ around x_∞ and $f_1(x_\infty) = f_2(x_\infty)$. Then for $i = 1$ and 2, there are elements g_i of G such that $g_i = \phi_N \circ f_i \circ \phi_M^{-1}$, where this equality holds in a set U where the maps are defined. For i sufficiently large, x_i belongs to $\phi_M^{-1}(U)$. In some neighbourhood of x_i, f_1 and f_2 agree. So g_1 and g_2 agree on some open set. Therefore, $g_1 = g_2$, since G is rigid. So, $f_1 = f_2$ on $\phi_M^{-1}(U)$. Therefore x_∞ is in V. □

Let M be a (G, X)-manifold, where G is rigid. Let $p: \tilde{M} \to M$ be the universal cover. Give $\tilde{M}$ its inherited (G, X)-structure. Then each covering transformation $\tau: \tilde{M} \to \tilde{M}$ is a (G, X)-isomorphism. Pick a basepoint x_0 in $\tilde{M}$ and a connected chart $\phi: U_0 \to X$ around it. Let $D: \tilde{M} \to X$ be the associated developing map. Note that $\phi \circ \tau: \tau^{-1}(U_0) \to X$ is a chart around $\tau^{-1}(x_0)$. We may use this chart to define D in a neighbourhood of $\tau^{-1}(x_0)$. Then $D|_{\tau^{-1}(U_0)} = g_\tau^{-1} \circ \phi \circ \tau$, for some element g_τ of G. So, the following commutes on $\tau^{-1}(U_0)$:

$$\begin{array}{ccc} \tilde{M} & \xrightarrow{\tau} & \tilde{M} \\ \downarrow D & & \downarrow D \\ X & \xrightarrow{g_\tau} & X \end{array}$$

By Lemma 10.2, this diagram commutes on all of M. If we compose two covering transformations τ and σ, then (by pasting two commutative diagrams together) $g_{\sigma\tau} = g_\tau g_\sigma$. So, we have a group homomorphism

$$\eta\colon \pi_1(M) \to G$$

known as the *holonomy*. It depends on the choice of chart $\phi\colon U_0 \to X$ and of basepoint. Different choices lead to a holonomy which differs by conjugation by an element of G.

In the case where X is a simply-connected Riemannian manifold, G is a group of isometries and M is complete, then $D\colon \tilde{M} \to X$ is a (G, X)-isomorphism. This identifies $\tilde{M}$ with X. Hence, we have the following.

Proposition 10.3. *Let X be a simply-connected Riemannian manifold and let G be a group of isometries. Let M be a complete (G, X)-manifold. Then $\eta(\pi_1(M))$ is a group of covering transformations of X, and M is the quotient $X/\eta(\pi_1(M))$.*

Now consider the case where $X = \mathbb{H}^n$ and $G = \mathrm{Isom}(\mathbb{H}^n)$, and where M is complete. Then, η sends each non-trivial element of $\pi_1(M)$ to an isometry of $\mathbb{H}^n$ which fixes no point of $\mathbb{H}^n$.

Corollary 10.4. *The holonomy homomorphism η is injective and its image contains no elliptic isometries other than the identity.*

11. Topological properties of complete hyperbolic manifolds

Lemma 11.1. *$\mathbb{H}^n$ has infinite volume.*

Proof. The volume of U^n is

$$\int_{(x_1,\dots,x_n)\in U^n} \frac{1}{x_n^n} dx_1 \dots dx_n = \infty$$ □

Proposition 11.2. *A complete finite volume hyperbolic manifold M has infinite fundamental group.*

Proof. If not, the universal cover $p\colon \tilde{M} \to M$ would be finite-to-one. Then $\tilde{M}$ would have finite volume. But $\tilde{M}$ is isometric to $\mathbb{H}^n$. □

Corollary 11.3. *S^n does not admit a hyperbolic structure for $n > 1$.*

Recall Thurston's geometrisation theorem:

Theorem. [Thurston] *Let M be a compact orientable irreducible atoroidal 3-manifold-with-boundary, such that ∂M is a non-empty collection of tori. Then either $M - \partial M$ has a complete finite volume hyperbolic structure, or M is homeomorphic to one of the following exceptional cases:*

1. $S^1 \times [0,1] \times [0,1]$
2. $S^1 \times S^1 \times [0,1]$
3. *the space obtained by gluing the faces of a cube as follows: arrange the six faces into three opposing pairs; glue one pair, by translating one face onto the other; glue another pair, by translating one face onto the other and then rotating through π about the axis between the two faces.*

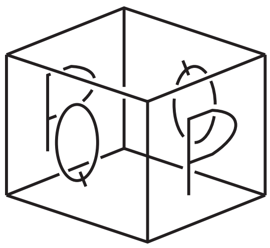

Figure 30.

We are now going to investigate to what extent these conditions are necessary.

Theorem 11.4. *Any complete hyperbolic 3-manifold M is irreducible.*

Proof. Let S^2 be a smoothly embedded 2-sphere in M. Since S^2 is simply-connected, the universal cover $p\colon \mathbb{H}^3 \to M$ restricts to a homeomorphism on each component of $p^{-1}(S^2)$. Pick one component S_1^2 of $p^{-1}(S^2)$. Then, by the Schoenflies theorem, the closure of one component of $\mathbb{H}^3 - S_1^2$ is homeomorphic to a closed 3-ball B_1 which is therefore compact. Therefore, there can only be finitely many covering translates of S_1^2 in B_1. Pick one innermost in B_1. This bounds a closed ball B_2. The covering translates of B_2 are all disjoint. Therefore, p projects B_2 homeomorphically to a closed ball in M. The boundary of this ball is the original S^2. □

We now investigate atoroidality. In fact any complete finite volume hyperbolic manifold is atoroidal, but the proof of this requires some understanding of the ends of these manifolds. Instead, we focus on the closed case. Even here, we need some extra geometric concepts.

Definition. Let M be a Riemannian manifold. For each $x \in M$, define the *injectivity radius* at x to be

$$\operatorname{inj}(x) = \sup\{\epsilon : \exp_x \text{ is injective on } B_\epsilon(0) \subset T_x M\}.$$

Note that there is such an $\epsilon > 0$, by Proposition 1.3 of the Introduction to Riemannian Manifolds.

Proposition 11.5. *Let M be a complete hyperbolic manifold. Let $p\colon \mathbb{H}^n \to M$ be the universal cover. For each $x \in M$, pick a point $\tilde{x} \in p^{-1}(x)$. Let*

$$i(\tilde{x}) = \sup\{\epsilon : \gamma(B_\epsilon(\tilde{x})) \cap B_\epsilon(\tilde{x}) = \emptyset \textit{ for all } \gamma \in \eta(\pi_1(M)) - \mathrm{id}\}.$$

Then $i(\tilde{x}) = \operatorname{inj}(x)$.

Proof.

Claim. $i(\tilde{x}) \geq \operatorname{inj}(x)$.

Pick $\epsilon < \operatorname{inj}(x)$ arbitrarily close to $\operatorname{inj}(x)$. Suppose that $\gamma(B_\epsilon(\tilde{x})) \cap B_\epsilon(\tilde{x}) \neq \emptyset$, for some $\gamma \in \eta(\pi_1(M)) - \mathrm{id}$. Let z be a point in their intersection, which we may assume to be on the geodesic joining $\tilde{x}$ to $\gamma(\tilde{x})$. Then $z = \exp_{\tilde{x}}(v_1) = \exp_{\gamma(\tilde{x})}(v_2)$ for vectors v_1 and v_2 with length less than ϵ. Note that if $(T\gamma)_{\tilde{x}}(v_1) = v_2$, then γ would preserve the geodesic between $\tilde{x}$ and $\gamma(\tilde{x})$ and would reverse its orientation. This would imply that γ fixed some point in $\mathbb{H}^n$, which is impossible. Hence, $(Tp)_{\tilde{x}}(v_1) \neq (Tp)_{\gamma(\tilde{x})}(v_2)$. These are distinct vectors in $T_x M$ with length less than ϵ which map to the same point $p(z)$ in M. This contradicts the definition of $\operatorname{inj}(x)$.

Claim. $\operatorname{inj}(x) \geq i(\tilde{x})$.

Pick $\epsilon < i(\tilde{x})$ arbitrarily close to $i(\tilde{x})$. Then $p|_{B_\epsilon(\tilde{x})}$ is an isometry onto its image. Hence, we may take its inverse ϕ to be a chart for M around x. Since $\exp_{\phi(x)}$ is injective, $\exp_x$ is injective on $B_\epsilon(0) \subset T_x M$. □

Corollary 11.6. *For a complete hyperbolic manifold M,* inj *is a continuous function on M.*

Proof. Let x_1 and x_2 be points in M. Let $\tilde{x}_1$ and $\tilde{x}_2$ be points in $p^{-1}(x_1)$ and $p^{-1}(x_2)$ such that $d(\tilde{x}_1, \tilde{x}_2) = d(x_1, x_2)$. Then, a ball of radius ϵ around $\tilde{x}_1$ contains a ball of radius $\epsilon - d(\tilde{x}_1, \tilde{x}_2)$ around $\tilde{x}_2$. So, by Proposition 11.5,

$$\text{inj}(x_2) \geq \text{inj}(x_1) - d(x_1, x_2).$$

The same is true with the rôles of x_1 and x_2 reversed. Hence,

$$|\text{inj}(x_1) - \text{inj}(x_2)| \leq d(x_1, x_2).$$

So, inj is continuous. □

Corollary 11.7. *If M is a closed hyperbolic manifold, there is a positive lower bound on* inj(x) *for all $x \in M$.*

Proposition 11.8. *Let M be a closed hyperbolic manifold. Then the image of each non-trivial element of $\pi_1(M)$ under η is a loxodromic isometry.*

Proof. Let γ be non-trivial element of $\pi_1(M)$. By Corollary 10.4, $\eta(\gamma)$ is non-trivial and non-elliptic. It is therefore parabolic or loxodromic. If it is parabolic, then by Proposition 4.5, it is conjugate to the isometry as in Example 2 of §2 with $\lambda = 1$. Recall that this is a Euclidean isometry $h\colon U^n \to U^n$ which fixes the n^{th} co-ordinate. For $(x_1, \dots, x_n) \in U^n$,

$$d_{hyp}(h(x_1, \dots, x_n), (x_1, \dots, x_n)) < d_{Eucl}(h(x_1, \dots, x_n), (x_1, \dots, x_n))/x_n \to 0$$

as $x_n \to \infty$. So, $\text{inj}(p(x_1, \dots, x_n)) \to 0$ as $x_n \to \infty$. This contradicts Corollary 11.7. □

Lemma 11.9. *Let f and g be commuting functions from a set to itself. Then g maps the fixed point set of f to itself.*

Proof. If $f(x) = x$, then $g(x) = gf(x) = fg(x)$. □

Theorem 11.10. *Let M be a closed hyperbolic n-manifold. Then no subgroup of $\pi_1(M)$ is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$.*

Proof. Let γ_1 and γ_2 be commuting elements of $\pi_1(M)$. Then, $\eta(\gamma_1)$ and $\eta(\gamma_2)$ are loxodromic, by Proposition 11.8. Since they commute, $\eta(\gamma_1)$ maps the fixed point set of $\eta(\gamma_2)$ to itself. It therefore preserves the geodesic α left invariant by $\eta(\gamma_2)$. It cannot reverse the orientation of α, since then it would have a fixed point on α.

Thus, $\eta(\gamma_1)$ has the same fixed point set as $\eta(\gamma_2)$. They both translate points on α some fixed hyperbolic distance along α. There is a uniform lower bound on this translation length for all non-identity elements of the group generated by $\eta(\gamma_1)$ and $\eta(\gamma_2)$. Therefore, some power of $\eta(\gamma_1)$ equals some power of $\eta(\gamma_2)$. So, γ_1 and γ_2 do not generate a subgroup isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$. □

Remark. The above theorem is false if 'closed' is replaced with 'complete and finite volume'. For example, it can be shown that for any knot K in S^3 other than the unknot, the map $i_*\colon \pi_1(\partial N(K)) \to \pi_1(S^3 - K)$ is injective, where i is the inclusion map of the boundary of the tubular neighbourhood $N(K)$ into $S^3 - K$. However, for 'most' knots K, $S^3 - K$ admits a complete finite volume hyperbolic structure. In these cases, $\eta(i_*\pi_1(\partial N(K)))$ is group of commuting parabolic isometries.

Corollary 11.11. *The n-manifold $S^1 \times \ldots \times S^1$ $(n > 1)$ does not admit a hyperbolic structure.*

Corollary 11.12. *A closed hyperbolic 3-manifold is atoroidal.*

Part III Hyperbolic Manifolds

Lent 1999

Examples Sheet 1

1. Show that the three interior angles of a triangle in the hyperbolic plane add up to less than π.

2. If α_1 and α_2 are disjoint geodesics in $\mathbb{H}^3$ which do not share a point at infinity, show that there is a geodesic α_3 which intersects both α_1 and α_2 orthogonally.

3. For an element $A \in PSL(2, \mathbb{C})$, let $\mathrm{tr}(A)$ denote its trace (which is defined up to sign). If $A \neq \pm\mathrm{id}$, show that the corresponding isometry of $\mathbb{H}^3$ is

(i) parabolic if $\mathrm{tr}(A) = \pm 2$

(ii) elliptic if $\mathrm{tr}(A) \in (-2, 2) \subset \mathbb{R}$,

(iii) loxodromic if $\mathrm{tr}(A) \in \mathbb{C} - [-2, 2]$.

In cases (ii) and (iii), show that $\mathrm{tr}(A)$ determines the conjugacy class of the isometry. In these cases, show how the angle of rotation and the hyperbolic translation length along the invariant geodesic can be calculated from $\mathrm{tr}(A)$.

4. Show that any two non-degenerate ideal triangles in $\mathbb{H}^2$ are isometric. Is the same true for ideal quadrilaterals? What about ideal tetrahedra in $\mathbb{H}^3$?

5. Show that, for *any* knot K in S^3, $S^3 - K$ admits an incomplete hyperbolic structure.

6. Let P be a non-degenerate hyperbolic polyhedron. Show that ∂P is the union of the facets of P and that P is the convex hull of its vertices.

7. Show that if P is non-degenerate polyhedron in $\mathbb{H}^3$ and V is the vertices of P, then $\partial P - V$ inherits an (incomplete) hyperbolic structure. Does this extend to a hyperbolic structure on all of ∂P? Show that if P' is a non-degenerate ideal polyhedron in $\mathbb{H}^3$, then $\partial P'$ inherits a complete hyperbolic structure.

8. Let P be a dodecahedron, namely the polyhedron with twelve pentagonal faces, shown overleaf.

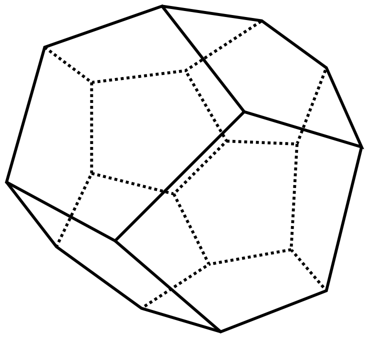

Let M be the space obtained by gluing each facet of P to the one opposite it, via a clockwise twist of $3\pi/5$. This is the Seifert-Weber dodecahedral space. Impose a hyperbolic structure on it. [This is necessarily complete and finite volume, and hence unique up to isometry, by Mostow Rigidity.]

9. Construct a complete hyperbolic structure on $S^1 \times \mathbb{R}^{n-1}$.

10. Show that the thrice-punctured 2-sphere S admits a complete hyperbolic structure, obtained by gluing two ideal triangles along their edges via isometries. [This is in fact the unique complete hyperbolic structure on S, up to isometry.] Show, however, that for 'most' ways of gluing the ideal triangles via isometries, the result is an incomplete hyperbolic structure on S.

11. [Hard] Generalise the technique for the construction of the hyperbolic structure on the figure-eight knot complement given in the lectures: construct a hyperbolic structure on the complements of the following links (which, if you do it correctly, will be complete and have finite volume):

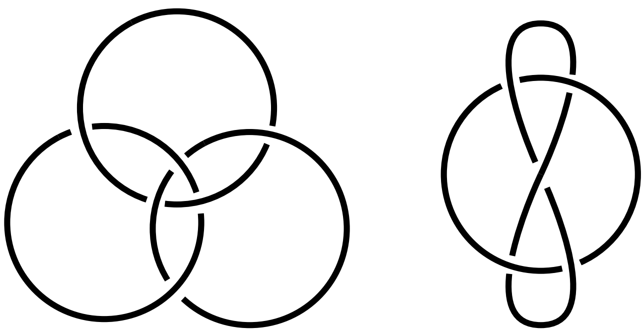

Part III Hyperbolic Manifolds

Lent 1999

Examples Sheet 2

1. Let $\tilde{M}$ be the universal cover of a Riemannian manifold M. Verify that a path $\alpha: I \to M$ is a geodesic if and only if some lift of α is a geodesic.

2. Show that if α is a closed geodesic in a complete hyperbolic manifold, then $\eta([\alpha])$ is loxodromic. Deduce that there are no simple closed geodesics in the hyperbolic structure on the thrice-punctured sphere given in Question 10 on Example Sheet 1. (A geodesic is *closed* if it factors through $\mathbb{R} \to S^1 \to M$. A non-closed geodesic $\alpha: \mathbb{R} \to M$ is *simple* if α is injective. A closed geodesic is *simple* if the associated map $S^1 \to M$ is injective.)

3. Show that each homotopically non-trivial closed curve α in a compact hyperbolic n-manifold is freely homotopic to a unique closed geodesic β. In the case $n = 2$, show that β is simple if α was simple. (A *free homotopy* between two closed curves $\alpha_0, \alpha_1: S^1 \to M$ is a homotopy $H: S^1 \times [0,1] \to M$ such that $H|_{S^1 \times \{t\}} = \alpha_t$, for $t = 0$ and 1. The word 'free' is used to emphasise that no basepoints are involved.)

4. Construct a simple non-closed geodesic on each compact orientable hyperbolic 2-manifold.

5. Define a *Euclidean n-manifold* to be a Riemannian manifold, each point of which has an open neighbourhood isometric to an open subset of $\mathbb{E}^n$, where $\mathbb{E}^n$ is $\mathbb{R}^n$ with the standard Euclidean metric. Adapt the techniques of the lectures to show that the universal cover of any complete Euclidean n-manifold is isometric to $\mathbb{E}^n$. A theorem of Bieberbach asserts that any group of isometric covering transformations for $\mathbb{E}^n$ contains a finite index subgroup consisting only of translations. Deduce that any compact Euclidean n-manifold is finitely covered by $S^1 \times \ldots \times S^1$. [One can also define a *spherical n-manifold* to be a Riemannian manifold locally modelled on S^n. Again, any complete spherical manifold has universal cover S^n. However, the techniques of the lectures do not immediately give this fact: where do they break down?]

6. If M is any open subset of $\mathbb{H}^n$ and $\tilde{M}$ is its universal cover, what are the

possible images for $D(\tilde{M})$, where D is a developing map for $\tilde{M}$? [Hint: prove and use the fact that a local isometry $h\colon N \to N'$ between connected Riemannian manifolds is determined by $h(x)$ and $(Th)_x$ for any $x \in N$.]

7. Recall the hyperbolic structure on the compact orientable surface F_k ($k > 1$) given in Theorem 3.2.2, obtained by gluing the facets of a hyperbolic $4k$-gon P. Show that, for a suitable choice of basepoint, a fundamental domain for F_k is P.

8. Show that neither $S^1 \times S^1 \times (0,1)$ nor $S^1 \times (0,1) \times (0,1)$ admits a complete finite volume hyperbolic structure. Show however that they both admit an uncountable number of non-isometric complete (infinite volume) hyperbolic structures.

9. Show that if M is any complete hyperbolic 3-manifold, then $\mathbb{Z} \oplus \mathbb{Z} \oplus \mathbb{Z}$ is not a subgroup of $\pi_1(M)$.

10. For sufficiently small $\epsilon > 0$, determine $\mathrm{inj}^{-1}((0,\epsilon])$ for the complete hyperbolic structure on the figure-eight knot complement given in the lectures. Your description should be both topological and geometric.

11. Let Γ be the set of elements

$$\pm \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{PSL}(2,\mathbb{R})$$

such that a and d are odd integers, and b and c are even integers. Verify that Γ forms a subgroup of $\mathrm{PSL}(2,\mathbb{R})$. Show that it is a group of isometric covering transformations and hence that $\mathbb{H}^2/\Gamma$ inherits a complete hyperbolic structure. Show that this is isometric to the hyperbolic structure on the thrice-punctured sphere S given in Question 10 on Example Sheet 1. [Hint: $\pi_1(S)$ is a free group on two generators. Show that (a suitable choice of) η sends these generators to

$$\pm \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \text{ and } \pm \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}.$$

Hence, $\eta(\pi_1(S))$ is a subgroup of Γ. So, S covers $\mathbb{H}^2/\Gamma$. Since S has finite volume, this is a finite cover. Now show that this cover must be the identity.]

12. Construct, for each compact orientable surface F, a non-identity homeomorphism $h\colon F \to F$ such that $h \circ h$ is the identity. Let M be the result of $F \times [0,1]$ after gluing $F \times \{0\}$ to $F \times \{1\}$ via h. Show that there is a cover $M \to M$ which

has finite index greater than one. Deduce that M has zero Gromov norm and hence does not admit a hyperbolic structure. Why is M not a counter-example to the conjecture of Thurston given before the start of Section 1?