11
Geometry of hyperbolic space

This chapter gives a short introduction to the geometry of hyperbolic space and to Kleinian groups. We begin in Sections 11.1–11.4 with hyperbolic space and its group of automorphisms. Section 11.5 treats elementary and non-elementary Kleinian groups. Section 11.6 defines the limit set of a Kleinian group and gives some of its basic properties.

Section 11.7 explores the *Jørgensen inequality*, a beautiful and rather mysterious inequality that exploits the discreteness of a Kleinian group in an essential way. This leads in Section 11.8 to a fairly easy proof of the *Margulis lemma*, which gives a complete description of the *thin* parts of a hyperbolic 3-manifold. It is a 3-dimensional analogue of the collaring theorem (Theorem 3.8.3) and extends the plumbing picture of hyperbolic surfaces given in Section 3.8. I find the proof of Jørgensen's inequality hard to motivate, so I also give in Appendix D2 a different proof of the Margulis lemma, which is cruder, but much more natural (and which generalizes to higher dimensions).

Thurston's hyperbolization theorems require not just thinking about individual Kleinian groups, but also thinking about their limits. There are two notions of "limit of a sequence of Kleinian groups", with surprisingly different properties: algebraic limits, which are relatively easy to understand, and geometric limits, which can be amazingly complicated. Algebraic limits are described in Section 11.9, geometric limits in Section 11.10.

Thurston's second hyperbolization theorem requires the Klein-Maskit combination theorems, which will eventually allow us to glue together hyperbolic manifolds. These theorems, proved in Section 11.11, also illustrate the power of 3-dimensional thinking about Kleinian groups; it is quite hard to see why the combined groups as defined are discrete if one sticks to the action on the Riemann sphere.

Section 11.12 proves the 3-dimensional analogue of the Poincaré polygon theorem (Theorem 3.9.5). I also discuss examples of Kleinian groups, including some that come from arithmetic: the *Bianchi groups* $PSL_2(\mathcal{O}_d)$, where $\mathcal{O}_d$ is the ring of integers in the imaginary quadratic field $\mathbb{Q}(\sqrt{-d})$. These groups are the natural generalization of the Fuchsian group $\mathrm{PSL}_2(\mathbb{Z})$; they, their congruence subgroups, and the associated modular forms are of such immense importance that it seems impossible to ignore them, even though they are not directly relevant to the topics of this book.

Finally, Section 11.13 defines *geometrically finite Kleinian groups*. These are the groups that are classified by points of appropriate Teichmüller spaces, leading in due course to the *skinning lemma* in Volume 4.

11.1 The hyperboloid model

It is possible to define n-dimensional hyperbolic geometry using a system of axioms more or less like the axioms of Euclidean geometry. I find it much more straightforward to define it as the geometry of a certain subset of $\mathbb{R}^{n+1}$ with an appropriate metric.

There are many models of n-dimensional hyperbolic space $\mathbb{H}^n$, just as there are many models of the hyperbolic plane; we will explore several in the next section. The most natural – if not the easiest to visualize – is the *hyperboloid model*, which imitates the definition of the unit sphere.

We use the symbol $\mathbb{H}^n$ to denote both the hyperboloid model and more generally the abstract hyperbolic space with all its geometry (metric, lines, planes, angles, etc.). In this volume we denote the hyperbolic plane by $\mathbb{H}^2$, not by $\mathbb{H}$, as in Volume 1. Similarly we write $\mathbf{H}^2$ and $\mathbf{D}^2$ rather than $\mathbf{H}$ and $\mathbf{D}$.

Definition 11.1.1 (n-dimensional hyperbolic space) *Hyperbolic n-dimensional space* $\mathbb{H}^n \subset \mathbb{R}^{n+1}$ is the subset given by

$$-x_0^2 + (x_1^2 + \cdots + x_n^2) = -1, \quad x_0 > 0, \tag{11.1.1}$$

with the Riemannian metric induced by the pseudo-metric

$$-dx_0^2 + dx_1^2 + \cdots + dx_n^2\,; \tag{11.1.2}$$

this metric is the *infinitesimal hyperbolic metric*.

Exercise 11.1.2 generalizes Exercise 2.4.2 in Volume 1.

Exercise 11.1.2 Check that $\mathbb{H}^n$ is a Riemannian manifold: the pseudo-metric given by formula 11.1.2 induces a Riemannian metric on $\mathbb{H}^n$, i.e., the quadratic form $-dx_0^2+dx_1^2+\cdots+dx_n^2$ is positive definite on the tangent spaces to $\mathbb{H}^n$. ◇

Exercise 11.1.3 1. Check that the map $t \mapsto \begin{pmatrix}\cosh t\\ \sinh t\end{pmatrix}$ is an isometric parametrization $\mathbb{R} \to \mathbb{H}^1$.

2. Show that if $x_1, \dots, x_n$ satisfy $x_1^2+\cdots+x_n^2 = 1$, then the map $\mathbb{R} \to \mathbb{H}^n$ given by $t \mapsto \begin{pmatrix}\cosh t\\ (\sinh t)x_1\\ \vdots\\ (\sinh t)x_n\end{pmatrix}$ is an isometric inclusion. ◇

Notation 11.1.4 We denote by $E^{n,1}$ the space $\mathbb{R}^{n+1}$ with pseudo-metric

$$-dx_0^2 + (dx_1^2 + \cdots + dx_n^2). \tag{11.1.3}$$

The corresponding *pseudo inner product* is given by

$$\langle \mathbf{v}, \mathbf{w} \rangle := -v_0 w_0 + (v_1 w_1 + \cdots + v_n w_n). \tag{11.1.4}$$

In this notation, $\mathbb{H}^n$ is the subset of $E^{n,1}$ given by

$$\mathbb{H}^n = \left\{ \mathbf{v} \in E^{n,1} \mid \langle \mathbf{v}, \mathbf{v} \rangle = -1 \quad \text{and} \quad v_0 > 0 \right\}. \tag{11.1.5}$$

Definition 11.1.5 (Lines, planes, . . . in $\mathbb{H}^n$) In hyperbolic space $\mathbb{H}^n$, lines, planes, or, more generally, k-dimensional subspaces are intersections $V \cap \mathbb{H}^n$, where V is a $(k+1)$-dimensional vector subspace of $E^{n,1}$ such that $V \cap \mathbb{H}^n \neq \emptyset$.

Remark In Chapter 2, I consistently referred to "geodesics" rather than "lines"; now I will usually speak of lines and planes. The word "geodesic" encourages thinking extrinsically: putting oneself in the ambient $E^{n,1}$ and looking at $\mathbb{H}^n$ "from the outside". I hope that speaking of lines and planes in $\mathbb{H}^n$ encourages a different way of thinking: living in $\mathbb{H}^n$ and seeing what *straight* means intrinsically. Why encourage a change of attitude now? For one thing, an ability to take both points of view is good in itself. But besides, for $\mathbb{H}^2$ the ambient $E^{2,1}$ is our ordinary space, with a "funny metric"; it is fairly easy to imagine and draw. But $E^{3,1}$, the ambient space of $\mathbb{H}^3$, is 4-dimensional and much harder to imagine. It is usually easier to imagine living in $\mathbb{H}^3$ than looking at $\mathbb{H}^3$ from $\mathbb{R}^4$. This approach has been championed by Thurston; not everyone can aspire to his skill, but I think everyone can gain from the attempt. △

Remark The space $E^{n,1}$ may seem like a slightly pathological curiosity: who cares about spaces with pseudo inner products? As briefly discussed in Section 2.4, nothing could be further from the truth: when $E^{3,1}$ is given the *Lorentz metric*

$$-c^2\, dt^2 + dx^2 + dy^2 + dz^2, \tag{11.1.6}$$

(with c the speed of light), then $E^{3,1}$ is *Minkowski space*, representing spacetime. (Since the units of c are distance/time, all four terms have units distance squared and it makes sense to add them.) Physicists, especially those who specialize in relativity, systematically set $c := 1$; then Minkowski space is exactly $E^{3,1}$.

Minkowski space and the Lorentz metric are the setting for the special theory of relativity; they are also the local model for general relativity. As such, they form the framework for much of modern physics, in particular electromagnetism and gravitation; physicists wish that quantum theory would fit into this framework too.

In the language of special relativity, points of $E^{3,1}$ are called *events*. A vector $\mathbf{v}$ connecting events is *time-like* if $\langle \mathbf{v}, \mathbf{v} \rangle < 0$, *space-like* if $\langle \mathbf{v}, \mathbf{v} \rangle > 0$. A time-like vector $\mathbf{v}$ can point to the future (if its coordinate v_0 is positive) or to the past (if v_0 is negative). Signals can travel from one event $\mathbf{p}_1$ to another event $\mathbf{p}_2$ if $\mathbf{p}_2 - \mathbf{p}_1$ is a time-like vector pointing to the future.

The vectors $\mathbf{v}$ such that $\langle \mathbf{v}, \mathbf{v} \rangle = 0$ are called *light-like*; the set of light-like vectors is called the *light cone* C; the *positive light cone* C^+ consists of light-like vectors pointing to the future, i.e., with $v_0 > 0$. A light-like vector can point to the future or to the past, and a *light ray* can travel from event $\mathbf{p}_1$ to event $\mathbf{p}_2$ if $\mathbf{p}_2 - \mathbf{p}_1$ is a light-like vector pointing to the future.

There is an immense literature on Minkowski space. Einstein's fundamental question about events is: can they be connected by a light ray? Essentially he says that the geometry of spacetime is completely controlled by the answer to that question. △

We will use the above language (*time-like*, *space-like*, *light cone*), inspired by physics, in $E^{n,1}$ for all n.

Automorphisms of $\mathbb{H}^n$

The notation Aut, like Hom, depends on context: $\operatorname{Aut} X$ is the group of automorphisms $X \to X$: maps that preserve the structure of X. If X is a metric space (for instance $\mathbb{H}^n$), $\operatorname{Aut} X$ consists of isometries; if X is a Riemann surface (for instance $\mathbb{P}^1$), $\operatorname{Aut} X$ consists of analytic isomorphisms; if X is a vector space, $\operatorname{Aut} X$ consists of linear isomorphisms. Theorem 1.8.2 in Volume 1 classified the automorphisms of simply connected Riemann surfaces.

In this book we are interested in orientation-preserving automorphisms of $\mathbb{H}^n$, which we denote by $\operatorname{Aut} \mathbb{H}^n$ (we won't be interested in orientation-reversing isometries, although in other settings, for instance, classifying nonorientable manifolds, they are important).

Remark 11.1.6 We will see in this section that $\operatorname{Aut} \mathbb{H}^n$ is identical to $\mathrm{SO}^+(E^{n,1}) \subset \mathrm{SL}_{n+1} \mathbb{R}$, the set of $(n+1) \times (n+1)$ real matrices that preserve the quadratic form $-x_0^2 + x_1^2 + \cdots + x_n^2$ and preserve the positive light cone and have determinant 1.

This is usually written $\operatorname{Aut} \mathbb{H}^n = \mathrm{SO}^+(n, 1)$. Viewing an automorphism as an element of $\operatorname{Aut} \mathbb{H}^n$ or as an element of $\mathrm{SO}^+(n, 1)$ is just a choice of