

Non-euclidean virtual reality II: explorations of $\mathbb{H}^2 \times \mathbb{E}$

Vi Hart Andrea Hawksley
eleVR eleVR
HARC HARC

Elisabetta A. Matsumoto
School of Physics
Georgia Institute of Technology

Henry Segerman
Department of Mathematics
Oklahoma State University

Abstract

We describe our initial explorations in simulating non-euclidean geometries in virtual reality. Our simulation of the product of two-dimensional hyperbolic space with one-dimensional euclidean space is available at h2xe.hypernom.com.¹




Figure 1: A view from $\mathbb{H}^2 \times \mathbb{E}$.

We are used to living in three-dimensional euclidean space, and our day-to-day experiences of curvature centre around surfaces embedded in $\mathbb{E}^3$. In the study of topology, the closed two-dimensional surfaces are the sphere, the torus, the two-holed torus, the three-holed torus, and so on. Thinking of these surfaces topologically, they don't come with a particular choice of *geometry* – that is, we can think about a surface as if it were made from plasticine – without knowing what lengths and angles mean on the surface. There are however particularly nice geometries for these surfaces: they are *homogeneous*, meaning that the geometry is the same everywhere in the space, and *isotropic*, meaning the geometry is the same in every direction. We have spherical geometry for the sphere, euclidean geometry for the torus, and hyperbolic geometry for all of the others. In three dimensions, the story is more complicated. Thurston's geometrization conjecture, proved by Perelman [3], gives eight geometries that a three-manifold can take (although the manifold may need to be decomposed into pieces, each with one of the eight geometries). The eight geometries are S^3 , $\mathbb{E}^3$, $\mathbb{H}^3$,

¹The code is available at github.com/henryseg/H2xE_VR.




Figure 2: Views of the $\{4, 6\}$ tiling of $\mathbb{H}^2 \times \mathbb{E}$. We draw the honeycomb to a depth of seven steps from the central cube.

$S^2 \times \mathbb{E}$, $\mathbb{H}^2 \times \mathbb{E}$, Nil , $\widetilde{\text{PSL}_2(\mathbb{R})}$ and Solv [4]. The first three are again homogeneous and isotropic: spherical, three-dimensional euclidean and hyperbolic geometries. The other five are homogeneous, but not isotropic: the geometries look different when you look in different directions. The geometries $S^2 \times \mathbb{E}$ and $\mathbb{H}^2 \times \mathbb{E}$ are mixtures of the two-dimensional geometries and one-dimensional euclidean space. The last three are more complicated versions of these mixed geometries: Nil and $\widetilde{\text{PSL}_2(\mathbb{R})}$ are “twisted” versions of $\mathbb{E}^2 \times \mathbb{E}$ and $\mathbb{H}^2 \times \mathbb{E}$ respectively, whilst Solv can be seen as a different kind of product of $\mathbb{E}^2$ and $\mathbb{E}$, in which moving in the $\mathbb{E}$ direction stretches and squeezes the $\mathbb{E}^2$ in two orthogonal directions.

Jeff Weeks’ software *Curved Spaces* [5] is a “flight simulator for multiconnected universes”. It simulates movement within a selection of closed three-dimensional manifolds, with S^3 , $\mathbb{E}^3$ and $\mathbb{H}^3$ geometries. Each of these is viewed as if we are living inside the space and seeing objects in the space via rays of light that travel along *geodesics* in the space. That is, light travels along paths of shortest distance.

A space such as $\mathbb{H}^3$ is appealing to create and explore because like euclidean space, $\mathbb{E}^3$, it is homogeneous and isotropic. Yet, when you move through the space, you encounter several unexpected consequences, including the effects of *holonomy* and *geodesic deviation*. Holonomy is a rotation of your reference frame as you traverse a loop in a curved space, which leads to the virtual floor becoming tilted away from horizontal. Geodesic deviation is the divergence or convergence of two geodesics that start out in parallel directions, which leads to the floor falling out from beneath your feet as you walk forward [1].

We sought a gentler introduction to hyperbolic geometry. One of the other eight Thurston geometries, $\mathbb{H}^2 \times \mathbb{E}$ – the cartesian product of the hyperbolic plane and the euclidean line – enables the user to traverse the hyperbolic plane horizontally as they walk through a room, yet it retains familiar euclidean geometry in the vertical direction. In this space, the virtual floor remains in the same place as the real-life floor. Jeff Weeks

has also explored $\mathbb{H}^2 \times \mathbb{E}$ via a similar method to ours in unpublished work.

There are four ingredients that go into our virtual reality simulation of $\mathbb{H}^2 \times \mathbb{E}$ as outlined in this paper:

1. A way to describe the points of $\mathbb{H}^2 \times \mathbb{E}$ numerically, i.e. a *model* of $\mathbb{H}^2 \times \mathbb{E}$,
2. A way to convert points in the model into points in $\mathbb{E}^3$ that we can then draw on screen,
3. A way to move around $\mathbb{H}^2 \times \mathbb{E}$ using the motion inputs from the virtual reality headset, and
4. A set of landmarks in $\mathbb{H}^2 \times \mathbb{E}$ to draw, to help the viewer navigate the space – we use a tiling of $\mathbb{H}^2 \times \mathbb{E}$.

1 The Model of $\mathbb{H}^2 \times \mathbb{E}$

The space $\mathbb{H}^2 \times \mathbb{E}$ is the cartesian product of the two spaces $\mathbb{H}^2$ and $\mathbb{E}$. It is homogeneous, but unlike $\mathbb{E}^3$, S^3 and $\mathbb{H}^3$, $\mathbb{H}^2 \times \mathbb{E}$ is not isotropic: the geometry is different when we look along an $\mathbb{H}^2$ versus an $\mathbb{E}$ direction. Our model is the cartesian product of the hyperboloid model of $\mathbb{H}^2$ with the real line. The model of $\mathbb{H}^2 \times \mathbb{E}$ lives in four-dimensional Minkowski space $\mathbb{E}^{3,1}$ as the set of points $\{(x, y, z, w) \in \mathbb{E}^{3,1} \mid x^2 + y^2 = w^2 - 1, w > 0\}$. The coordinates x, y and w are used for the hyperboloid model of $\mathbb{H}^2$, while the coordinate z is used for $\mathbb{E}$. We use the explicit parametrisation given by the map $\phi : \mathbb{R}_{\geq 0} \times [0, 2\pi) \times \mathbb{R} \rightarrow \mathbb{E}^{3,1}$ given by

$$\phi(\rho, \theta, z) = (\sinh \rho \cos \theta, \sinh \rho \sin \theta, z, \cosh \rho).$$

2 Drawing points in $\mathbb{H}^2 \times \mathbb{E}$ on the screen

To draw points on screen, we cannot use the same trick as we did in $\mathbb{H}^3$ of ignoring the inverse of the exponential map and using the Klein model instead [1]. Naïvely, we might try converting the x, y and w coordinates of a point into a point in the Klein model of $\mathbb{H}^2$, sitting inside of the unit disk in $\mathbb{E}^2$, and draw points in the cartesian product of the Klein model and $\mathbb{E}$, using z for the value in $\mathbb{E}$. If we did this then we would see objects in the space, but if we tried to move towards an object we would miss. This is because the light rays in $\mathbb{E}^3$ which connect an object to our eyes do not generally lie along the geodesics of $(\text{Klein model}) \times \mathbb{E}$. To properly account for this, we shall digress to a discussion of geodesics and the exponential map in $\mathbb{H}^2 \times \mathbb{E}$.

Whilst geodesics in the Klein model are straight lines in the euclidean space in which the model lives, geodesics in $(\text{Klein model}) \times \mathbb{E}$ have an altogether different flavour due to the anisotropic nature of $\mathbb{H}^2 \times \mathbb{E}$. If we are at the origin, looking in the diagonal direction shown in Figure 3, we see objects that are on the red line. But when we move in that direction we will arrive at objects that are on the blue geodesic.

So, we need to return to the original plan in our previous paper [1] of calculating the inverse of the exponential map in order to draw objects correctly. To do this, we need to work out parametrisations of geodesics in our model of $\mathbb{H}^2 \times \mathbb{E}$ in $\mathbb{E}^{3,1}$. Following the calculations outlined in Appendix A1, they have the form

$$\gamma(t) = \{\sinh(\rho_0 t) \cos \theta_0, \sinh(\rho_0 t) \sin \theta_0, z_0 t, \cosh(\rho_0 t)\}.$$

To draw points on screen, we construct the inverse of the exponential map. Given a point p in $\mathbb{H}^2 \times \mathbb{E}$, this tells us what initial velocity in the tangent space we need in order to reach p in $\mathbb{H}^2 \times \mathbb{E}$ at time $t = 1$. The inverse of the exponential map sends points on the $(\text{hyperboloid}) \times \mathbb{E}$ in $\mathbb{E}^{3,1}$ to points in $\mathbb{E}^3$ via

$$(x, y, z, w) \mapsto \left(\frac{\operatorname{arcsinh}(\sqrt{x^2 + y^2})}{\sqrt{x^2 + y^2}} x, \frac{\operatorname{arcsinh}(\sqrt{x^2 + y^2})}{\sqrt{x^2 + y^2}} y, z \right).$$






Figure 3: A geodesic in the $(\text{Klein model}) \times \mathbb{E}$ (blue), and a straight line in $\mathbb{E}^3$ (red).



(a) Our colouring on the $\{4, 6\}$ tiling in the Poincaré disk model.



(b) A fundamental domain for the tiling, drawn in the Poincaré disk model.



(c) The fundamental domain glued up to form a genus two surface.

Figure 4: Our colouring of the $\{4, 6\}$ tiling.

The images of geodesics in $\mathbb{H}^2 \times \mathbb{E}$ are straight lines in the tangent space. Thus if you look at a point in the virtual reality experience and move towards it, you will eventually reach it.

3 Moving through $\mathbb{H}^2 \times \mathbb{E}$

Once we know the relationship between the viewer's motions and the headset output, we must convert them to motions in $\mathbb{H}^2 \times \mathbb{E}$. Moving through this space now consists of translating by isometries inherited from $\mathbb{E}^{3,1}$ that leave the model unchanged. Turning this set of isometries into the exponential map (in the Lie algebra sense, see [1] for further details) can be handled in a variety of ways. As both $\mathbb{H}^2$ and $\mathbb{E}$ can each be embedded in a flat space of one higher dimension, the exponential map could be encoded in a block diagonal 5×5 matrix, but graphics cards are not optimised for this type of matrix multiplication. Instead, we can take advantage of $\mathbb{H}^2 \times \mathbb{E}$ being a product space to realize motions of the $\mathbb{H}^2$ factor using a 3×3 matrix (implemented in exactly the same way as we did for $\mathbb{H}^3$ [1]) while independently encoding translation in the $\mathbb{E}$ factor via simple addition.

4 The $\{4, 6\}$ tiling of $\mathbb{H}^2$ and a colouring of $\mathbb{H}^2 \times \mathbb{E}$

Eventually, we plan to put recognisable, human scale objects in our simulation of $\mathbb{H}^2 \times \mathbb{E}$, for people to interact with. For now, we pattern it with cubes, as we did for $\mathbb{H}^3$ [1]. In each horizontal $\mathbb{H}^2$ slice, we draw the $\{4, 6\}$ Schläfli symbol tiling: squares (4 sides), with 6 meeting at each vertex. These horizontal squares are cross-sections of our cubes. At regular intervals in the $\mathbb{E}$ direction, one cube ends, and the next begins. Because this is a product space, we have freedom to choose the height of the cubes in $\mathbb{E}$ relative to the side length of the $\{4, 6\}$ tiling of $\mathbb{H}^2$. We choose this height so that when viewed from its centre, a cube in $\mathbb{H}^2 \times \mathbb{E}$ has its vertices coincide with the vertices of a euclidean cube. The inverse of the exponential map associates each of these vertices to a point in the tangent space that is a distance $\text{arcsinh}(\sqrt{2})/\sqrt{2}$ away from the centre. We take the vertices in the meshes that make up the cubes and their decorations in $\mathbb{E}^3$ and move them upstairs into $(\text{hyperboloid}) \times \mathbb{E}$ via the map from the Klein model to the hyperboloid model according to the coordinate transformation $(x, y, z) \mapsto (x/\sqrt{1-x^2-y^2}, y/\sqrt{1-x^2-y^2}, z, 1/\sqrt{1-x^2-y^2})$. Once in the hyperboloid model, transformations are handled through the exponential map. Unlike $\mathbb{H}^3$ which is isotropic, the anisotropy of $\mathbb{H}^2 \times \mathbb{E}$ causes the straight lines bounding the faces of the cube in $\mathbb{E}^3$ to become curved as you look along a diagonal direction in $\mathbb{H}^2 \times \mathbb{E}$.

In Figure 2, we colour each layer of cubes in essentially the same way, coming from a colouring of the




Figure 5: The aspect ratio of the cubes scales differently in the $\mathbb{H}^2$ and $\mathbb{E}$ directions.

$\{4, 6\}$ tiling. Subsequent layers have their colours “rotated” slightly, around the circle of colours from black to red to white to cyan. As in our colouring of the $\{4, 3, 6\}$ honeycomb [1], the colouring on each layer comes from lifting a colouring of a closed manifold, in this case the genus two surface (see Figure 4).

5 Revenge of curved spaces in virtual reality

As with $\mathbb{H}^3$ [1], our experience in $\mathbb{H}^2 \times \mathbb{E}$ differs viscerally from in $\mathbb{E}^3$. One of the most striking differences stems from the lack of isotropy in $\mathbb{H}^2 \times \mathbb{E}$. As we move forwards, the aspect ratio of objects in front of us appears to change. Figure 5 shows a sequence of views, moving forwards from one cube into the next. The octagon in front of us starts out tall and thin, and becomes wider as we approach. The octagons further into the distance are even thinner. The reason for this behaviour in $\mathbb{H}^2 \times \mathbb{E}$ is that an object’s apparent height goes down linearly relative to its distance from us (just as it does in $\mathbb{E}^3$) but its apparent width scales exponentially. This doesn’t happen in $\mathbb{H}^3$ – both width and height scale exponentially, as $\mathbb{H}^3$ is isotropic. Thus, the perceived aspect ratio of objects does not change.

An interesting feature of $\mathbb{H}^2 \times \mathbb{E}$ being a product space is that parallel transport affects directions differently. In $\mathbb{H}^3$ you experience holonomy – when you traverse a loop in any plane you experience the world as having rotated. This is because the vector pointing “up” gets transported around the loop and comes back as rotated [1]. In $\mathbb{H}^2 \times \mathbb{E}$, you experience this phenomenon when looking along the $\mathbb{E}$ direction – that is to say, the loop lies in the $\mathbb{H}^2$ plane (see Figure 6). The frame is transported in the same it was in any direction in $\mathbb{H}^3$. However, if you look directly along the $\mathbb{H}^2$ plane and traverse a loop, first moving to the right for distance d along the $\mathbb{H}^2$ direction, then up along $\mathbb{E}$ for distance d , then to the left along $\mathbb{H}^2$ and finally returning to the starting point by travelling down along $\mathbb{E}$, you will notice that the world has not rotated. This is because pure isometries in $\mathbb{H}^2$ and $\mathbb{E}$ commute as $\mathbb{H}^2 \times \mathbb{E}$ is a cartesian product.

6 The In-Space View: Visualising $\mathbb{H}^2 \times \mathbb{E}$ and Parallax

In the virtual reality headset, a separate image is rendered for each eye, creating the illusion of a three-dimensional environment. Humans judge the distance of objects from them using a number of different





Figure 6: Holonomy rotates reference frames after moving in a loop in curved space. Rotations in the $\mathbb{H}^2$ plane (column 2) have the same effects as looking in any direction in $\mathbb{H}^3$ (column 1). However other directions behave differently. (In hyperbolic units the distance from the centre of a cube to the centre of a face is $\text{arccosh}(\sqrt{3/2})$.)

techniques. One such technique is our use of *stereopsis*, the perception of depth inferred from differences in the visual input coming into our two eyes. This employs *parallax*, the difference between the apparent position of an object viewed from different viewpoints. An object that is directly in front of us and close to us, will appear on the right side of the field of view of one's left eye, and on the left side of the field of view of one's right eye. An object that is directly in front of us but very far away from us will appear in the centre of the field of views from each eye, assuming that our eyes are both pointing directly forward (see Figure 7).

In particular, an inhabitant of euclidean space expects to see an object that is infinitely far away along



(a) Geodesics (light rays) from two eyes looking straight ahead would diverge in $\mathbb{H}^2$. (b) Eyes must point inward to see an object that is infinitely far away in $\mathbb{H}^2$.

Figure 7: Geodesic deviation is responsible for the phenomena associated with parallax in $\mathbb{H}^2$. See also [6].

rays of light that enter their two eyes along parallel geodesics. In negatively curved spaces however, geodesics that enter one's eyes along parallel directions diverge as we follow the light rays backwards – there can be no single object that those light rays came from. Instead, a euclidean visitor to hyperbolic space has to turn their eyes inwards to see an object that is infinitely far away. The effect is that all objects appear to be relatively close by. Objects in hyperbolic space are further away than they appear.

Euclidean visual effects create the illusion of depth by using the average pupil separation to create separate images from “cameras” centred at each eye. Where should we put the viewer's eyes in our simulations? This is a subtle point, to which we give two answers, both justifiable:

1. We could display the world as a hyperbolic being would see it – take the eyes to be virtual cameras placed inside of the curved space, each with its own inverse of the exponential map into its own distinct tangent space. Light rays follow geodesics that leave each point in the space and travel to the centre of each camera.
2. Instead, our present simulation uses one tangent space (centred at the headset) as the copy of $\mathbb{E}^3$ in which the euclidean user lives. The cameras are placed in that tangent space according to the user's physiology. In this case, the visually and physically perceived positions of objects on the user's body – for example a hand-held controller – are identical.

This being said, the user is unlikely to be able to distinguish between renderings in these two situations. The average interpupillary distance (54-68 mm in adult humans) is much smaller than the radius of curvature we use, 1 m, so differences will be minor.

7 Future directions: tracking objects in space and other geometries

Floating point errors provide a challenge in tracking objects in hyperbolic space, since the elements of our isometry matrices are exponential functions of the distance travelled in applying them. In our simulations so far, the user never actually leaves the central cube of the tiling: as they attempt to they are teleported from one side of it to the other, and the colours of the cubes are updated appropriately. We would not be able to use this trick for objects left in the world. One solution would be to record the location of an object by its position relative to a cube of our tiling, together with the cube's location recorded as a word in the tiling's generators. However, with this kind of data structure it is non-trivial to calculate, for example, which objects are close to your current position and so should be drawn on screen.

In addition to putting more recognisable objects and architecture into our simulations and allowing users to interact with objects, we would like to build similar simulations for the other Thurston geometries. Each of these geometries presents a unique challenge. Both S^3 and $S^2 \times \mathbb{E}$ have a multivalued exponential maps, thus we need to calculate the exponential map both in front of the viewer and behind them in order to draw

a complete image on the screen. Nil, Solv and $\widetilde{\text{PSL}_2(\mathbb{R})}$ don't have the ubiquitous standard models that spherical and hyperbolic space have. In our future implementations of Thurston geometries, we will use models of Emil Molnár [2] to create the exponential map as well as the set of isometries.

References

- [1] Vi Hart, Andrea Hawksley, Elisabetta A. Matsumoto, and Henry Segerman. Non-euclidean virtual reality I: explorations of $\mathbb{H}^3$. In *Proc. Bridges 2017*. Tessellations Publishing, 2017.
- [2] Emil Molnár. The projective interpretation of the eight 3-dimensional homogeneous geometries. *Beitrage zur Algebra und Geometrie (Contributions to Algebra and Geometry)*, 38(2):261–288, 1997.
- [3] Grisha Perelman. The entropy formula for the Ricci flow and its geometric applications. arXiv:0211159.
- [4] William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton Univ. Press, 1997.
- [5] Jeff Weeks. Curved Spaces. a flight simulator for multiconnected universes, available from <http://www.geometrygames.org/CurvedSpaces/>.
- [6] Jeff Weeks. Real-time rendering in curved spaces. *IEEE Computer Graphics and Applications*, 22(6):90–99, 2002.

A1 Geodesics in $\mathbb{H}^2 \times \mathbb{E}$

In order to implement other curved spaces in virtual reality, we will need to calculate the inverse of the (riemannian geometry) exponential map. Here, we present a formal method using differential geometry. The first step in calculating geodesics is finding the components of the metric tensor. In our parametrisation $\phi(\rho, \theta, z) = (\sinh \rho \cos \theta, \sinh \rho \sin \theta, z, \cosh \rho)$, the components of the metric tensor g_{ij} are given by $g_{ij} = \langle \partial_i \phi, \partial_j \phi \rangle$, where $i, j \in \{\rho, \theta, z\}$ and $\langle \cdot, \cdot \rangle$ denotes the inner product in the ambient space, $\mathbb{E}^{3,1}$. The three non-zero components of the metric tensor are $g_{\rho\rho} = g_{zz} = 1$ and $g_{\theta\theta} = (\sinh \rho)^{-2}$. Each of the derivatives $\partial_i \phi$ is a vector in the tangent space of $\mathbb{H}^2 \times \mathbb{E}$.

The derivative of a vector in the tangent space tells us the rate of change of that vector (both magnitude and direction) as we move along another vector in the tangent space. If the derivative has no component in any direction other than parallel to itself, then parallel transport in the direction of the vector preserves it. This is our condition for being a geodesic. Thus the geodesic $\gamma(t) = \phi(\rho(t), \theta(t), z(t))$ satisfies the equation $\nabla_{\dot{\gamma}} \dot{\gamma} = 0$, where $\dot{\gamma} = \frac{d\gamma}{dt}$ and ∇_X is the Levi-Civita connection. In coordinate form this may be re-written as

$$\partial_t^2 \gamma^\lambda + \sum_{\mu, \nu} \Gamma_{\mu\nu}^\lambda \partial_t \gamma^\mu \partial_t \gamma^\nu = 0,$$

where Christoffel symbols of the second kind $\Gamma_{\mu\nu}^\lambda = \frac{1}{2} \sum_\sigma g_{\lambda\sigma}^{-1} (\partial_\nu g_{\sigma\mu} + \partial_\mu g_{\sigma\nu} - \partial_\sigma g_{\mu\nu})$ are the components of the Levi-Civita connection and $\lambda, \mu, \nu, \sigma \in \{\rho, \theta, z\}$. In $\mathbb{H}^2 \times \mathbb{E}$, the only non-zero Christoffel symbols are $\Gamma_{\theta\theta}^\rho = -\cosh \rho \sinh \rho$, and $\Gamma_{\rho\theta}^\theta = \Gamma_{\theta\rho}^\theta = \coth \rho$. This set of coupled differential equations gives a formula for parametrised curves in the coordinate map $\gamma(t) = (\rho(t), \theta(t), z(t))$. In $\mathbb{H}^2 \times \mathbb{E}$ geodesics emanating from the origin $(0, 0, 0, 1)$ of the hyperboloid must satisfy the differential equations for curves $\rho''(t) = z''(t) = \theta'(t) = 0$, with boundary conditions $\rho(0) = 0$, $z(0) = 0$, $\theta(0) = \theta_0$, $\rho(1) = \rho_0$ and $z(1) = z_0$. Thus they are parametric curves in $\mathbb{E}^{3,1}$ of the form

$$\gamma(t) = \{\sinh(\rho_0 t) \cos \theta_0, \sinh(\rho_0 t) \sin \theta_0, z_0 t, \cosh(\rho_0 t)\}.$$