Ultrahyperbolic Neural Networks

Marc T. Law

NVIDIA

Abstract

Riemannian space forms, such as the Euclidean space, sphere and hyperbolic space, are popular and powerful representation spaces in machine learning. For instance, hyperbolic geometry is appropriate to represent graphs without cycles and has been used to extend Graph Neural Networks. Recently, some pseudo-Riemannian space forms that generalize both hyperbolic and spherical geometries have been exploited to learn a specific type of nonparametric embedding called ultrahyperbolic. The lack of geodesic between every pair of ultrahyperbolic points makes the task of learning parametric models (e.g., neural networks) difficult. This paper introduces a method to learn parametric models in ultrahyperbolic space. We experimentally show the relevance of our approach in the tasks of graph and node classification.

1 Introduction

Riemannian manifolds of constant curvature are the most common representation spaces in machine learning. They include the Euclidean space (of constant zero curvature), the d-sphere (of constant positive curvature) and the hyperbolic space (of constant negative curvature). The choice of a geometry to represent data mainly depends on the kind of relationship that needs to be described. For instance, Gromov [10] showed the relevance of hyperbolic geometry to represent trees (i.e., graphs without cycles). Since many hierarchies can be described as trees, hyperbolic representations have been used to represent hierarchical relationships (e.g., hypernymy between words [19]). Nonetheless, in many domains (e.g., social networks or protein structures), hierarchical graphs contain cycles.

In hyperbolic geometry, the considered manifold is not a vector space and is not equipped with the standard dot product. Therefore, most hyperbolic neural networks [5, 8, 18, 27] represent the weights of their last layer in the tangent space of some reference point. That tangent space is equipped with a positive definite metric tensor and the learned model can then be optimized with Riemannian gradient descent [1, 4]. In particular, since there exists a geodesic between any pair of points, the parameters are often optimized by using parallel transport (also called parallel translation) or the logarithm map. The Riemannian gradients are then parallel translated to the reference tangent space in which the model parameters lie. We refer the reader to [23] for a recent survey on hyperbolic neural networks.

Recently, Law & Stam [15] proposed *ultrahyperbolic* embeddings. They are a type of embedding that lies on a pseudo-Riemannian manifold of constant nonzero curvature [2, 21, 30]. Pseudo-Riemannian manifolds (also called semi-Riemannian manifolds) are generalizations of Riemannian manifolds where the nondegenerate metric tensor is not constrained to be positive definite [16]. In particular, when the metric tensor is not positive definite (e.g., when it is indefinite), the negative of the (pseudo-Riemannian) gradient is not a descent direction [9]. Law & Stam [15] proposed an efficient method to calculate a descent direction and learn ultrahyperbolic (nonparametric) embeddings. The main motivation of representing data on an ultrahyperbolic manifold is that it contains hyperbolic and spherical parts (see Fig. 1 and supp. material for details). It can then describe relationships specific to hyperbolic and spherical geometries (e.g., to represent parts of a graph that are trees or cycles) and is more flexible. Ultrahyperbolic embeddings were experimentally shown to be more appropriate than hyperbolic embeddings to represent hierarchical graphs with cycles on several datasets [15].

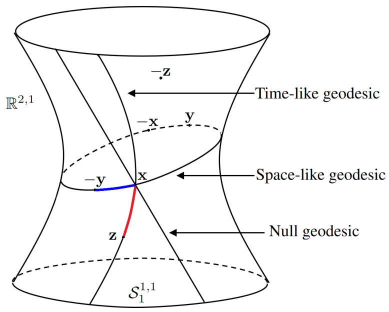

Figure 1: Geodesics of the pseudo-Riemannian quotient manifold $\mathcal{P}_1^{1,1} = \mathcal{S}_1^{1,1} / \pm 1$ embedded in $\mathbb{R}^{2,1}$. The point $[\mathbf{x}]$ of $\mathcal{P}_1^{1,1}$ is the pair $\{\mathbf{x}, -\mathbf{x}\}$. Any pair of points of $\mathcal{P}_1^{1,1}$ can be joined by a geodesic of $\mathcal{P}_1^{1,1}$. On the other hand, $\mathbf{x}$ and $-\mathbf{z}$ cannot be joined by an (unbroken) geodesic of $\mathcal{S}_1^{1,1}$. The length of the minimizing geodesic of $\mathcal{P}_1^{1,1}$ joining $[\mathbf{x}]$ and $[\mathbf{y}]$ is the length of the minimizing geodesic of $\mathcal{S}_1^{1,1}$ joining $\mathbf{x}$ and $-\mathbf{y}$ (in blue). The length of the geodesic of $\mathcal{P}_1^{1,1}$ joining $[\mathbf{x}]$ and $[\mathbf{z}]$ is the length of the geodesic of $\mathcal{S}_1^{1,1}$ joining $\mathbf{x}$ and $\mathbf{z}$ (in red). See details in the supp. material.

However, since there exist pairs of points on the ultrahyperbolic manifold considered in [15] that cannot be joined by an (unbroken) geodesic, gradients might not be parallel translated via a geodesic and the logarithm map joining two given points might not be defined. Directly extending hyperbolic neural networks [5, 8, 18, 27] to ultrahyperbolic space is then problematic.

In this paper, we propose a method to learn *ultrahyperbolic* representations with neural networks. Unlike [15], we consider the pseudo-Riemannian quotient manifold defined such that every point $\mathbf{x} = (x_0, \ldots, x_d)^\top$ is equivalent to its antipodal point $-\mathbf{x} = (-x_0, \ldots, -x_d)^\top$. In this way, for any other point $\mathbf{y}$, there always exists at least one geodesic joining $(\mathbf{x}, \mathbf{y})$ or $(-\mathbf{x}, \mathbf{y})$. We provide sufficient conditions to minimize a function defined on our quotient manifold. Since tangent vectors (hence gradients) of quotient manifolds are abstract objects, we explain how the function can be optimized with the *horizontal lift* operator. Our optimization framework is general, so we also introduce an extension to Graph Neural Networks (GNNs) [18] such that the activation representations at each layer of our GNN lie in ultrahyperbolic space. We then obtain a deep ultrahyperbolic model to represent graphs given as input. We evaluate our approach in different graph classification tasks.

2 Pseudo-sphere and Quotient Manifold

We extend the ultrahyperbolic manifold described in [15] (denoted by $\mathcal{S}_r^{p,q}$) to a quotient manifold denoted by $\mathcal{P}_r^{p,q}$ where (p, q) is the metric signature (see page 343 of [16]) of the pseudo-Riemannian manifold and $1/r^2$ is its curvature. The motivation is that any pair of points of $\mathcal{P}_r^{p,q}$ can be joined by at least one geodesic, which allows us to optimize parametric models. We consider three pseudo-Riemannian manifolds $\mathcal{P}_r^{p,q} \subset \mathcal{S}_r^{p,q} \subset \mathbb{R}^{p+1,q}$ that we define below. We explain how $\mathcal{P}_r^{p,q}$ generalizes elliptic and hyperbolic geometries in the special cases where $q = 0$ and $p = 0$, respectively.

Notation. We denote points on a smooth manifold $\mathcal{M}$ [16] by boldface Roman characters $\mathbf{x} \in \mathcal{M}$. $[\mathbf{x}] := \{\mathbf{x}, -\mathbf{x}\}$ is a pair of antipodal points. $T_\mathbf{x}\mathcal{M}$ is the tangent space of $\mathcal{M}$ at $\mathbf{x}$ and we write tangent vectors $\boldsymbol{\xi} \in T_\mathbf{x}\mathcal{M}$ in boldface Greek fonts. $\mathbb{R}^d$ is the d-dimensional Euclidean space equipped with the (positive definite) dot product $\langle \cdot, \cdot \rangle$ defined as $\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{x}^\top \mathbf{y}$. $\mathbf{I}$ is the identity matrix. The inverse function of the cosine (resp. hyperbolic cosine) is denoted by $\cos^{-1}$ (resp. $\cosh^{-1}$).

Ambient space $\mathbb{R}^{p+1,q}$. Our ambient space $\mathbb{R}^{p+1,q}$ is a vector space of dimensionality $d + 1 = p + q + 1 \in \mathbb{N}$ called *pseudo-Euclidean space* [21]. It is equipped with the following scalar product (i.e., nondegenerate symmetric bilinear form) of signature $(p + 1, q)$:

$$\forall \mathbf{x} = (x_0, \ldots, x_d)^\top, \ \mathbf{y} = (y_0, \ldots, y_d)^\top, \ \langle \mathbf{x}, \mathbf{y} \rangle_q := \sum_{i=0}^{p} x_i y_i - \sum_{j=p+1}^{d} x_j y_j = \mathbf{x}^\top \mathbf{G} \mathbf{y}, \tag{1}$$

where the signature matrix $\mathbf{G} = \mathbf{G}^{-1} = \mathbf{I}_{p+1,q}$ is the $(d+1) \times (d+1)$ diagonal matrix with the first $p+1$ diagonal elements equal to 1 and the remaining q equal to -1. Following general relativity literature and spacetime terminology [7], $\mathbb{R}^{p+1,q}$ has $p+1$ space dimensions and q time dimensions. Since it is a vector space, we can identify its tangent space to the space itself by means of the natural isomorphism $T_\mathbf{x}\mathbb{R}^{p+1,q} \approx \mathbb{R}^{p+1,q}$. Finally, the Euclidean space $\mathbb{R}^{d+1}$ is the special case of $\mathbb{R}^{d+1,0}$ which contains zero time dimension, and where $\mathbf{G} = \mathbf{I}_{d+1,0} = \mathbf{I}$.

Total space $\mathcal{S}_r^{p,q}$. Our total space $\mathcal{S}_r^{p,q}$ is a pseudo-sphere of radius $r > 0$ embedded in $\mathbb{R}^{p+1,q}$. It is the following hypersurface:

$$\mathcal{S}_r^{p,q} := \left\{ \mathbf{x} \in \mathbb{R}^{p+1,q} : \langle \mathbf{x}, \mathbf{x} \rangle_q = r^2 \right\}, \tag{2}$$

It is equivalent to work with the pseudo-hyperboloid $\mathcal{Q}_r^{q,p} := \{\mathbf{x} \in \mathbb{R}^{q,p+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{p+1} = -r^2\}$ and the pseudo-sphere $\mathcal{S}_r^{p,q}$ as they are anti-isometric to each other (see supp. material). Moreover, the radius $r > 0$ plays a role of scaling factor so we consider it to be 1 although it can be learned [5, 14]. Finally, both $\mathbf{x} \in \mathcal{S}_r^{p,q}$ and its antipodal point $-\mathbf{x}$ lie on $\mathcal{S}_r^{p,q}$ since $\langle \mathbf{x}, \mathbf{x} \rangle_q = \langle -\mathbf{x}, -\mathbf{x} \rangle_q$.

Quotient manifold $\mathcal{P}_r^{p,q}$. We consider as equivalence relation the two-element group $\{\pm 1\}$ consisting of the identity map $\mathbf{x} \mapsto \mathbf{x}$ and the antipodal map $\mathbf{x} \mapsto -\mathbf{x}$. This means that two points $\mathbf{x} \in \mathcal{S}_r^{p,q}$ and $\mathbf{y} \in \mathcal{S}_r^{p,q}$ are equivalent iff $\mathbf{y} = \mathbf{x}$ or $\mathbf{y} = -\mathbf{x}$. We define the following projective space:

$$\mathcal{P}_r^{p,q} := \mathcal{S}_r^{p,q} / \pm 1 = \mathcal{S}_r^{p,q} / \pm \mathbf{I} = \{\{\mathbf{x}, -\mathbf{x}\} : \mathbf{x} \in \mathcal{S}_r^{p,q}\}. \tag{3}$$

Every point of $\mathcal{P}_r^{p,q}$ is an unordered pair that we denote by $[\mathbf{x}] := \{\mathbf{x}, -\mathbf{x}\}$. Since $\mathcal{P}_r^{p,q}$ is a projective space, every point $[\mathbf{x}] \in \mathcal{P}_r^{p,q}$ can be interpreted as the intersection of the pseudo-sphere $\mathcal{S}_r^{p,q}$ with a line passing through the origin of $\mathbb{R}^{p+1,q}$. In some cases, it might be easier to interpret points of $\mathcal{P}_r^{p,q}$ as lines through the origin, and to study their properties when they intersect the pseudo-sphere. Each point $[\mathbf{x}] \in \mathcal{P}_r^{p,q}$ is also a submanifold of $\mathcal{S}_r^{p,q}$ and a discrete space.

In the following, we explain how $\mathcal{P}_r^{p,q}$ extends spherical geometry to elliptic geometry (i.e., when $q = 0$), or naturally describes the hyperboloid model of hyperbolic geometry (i.e., when $p = 0$).

Elliptic geometry ($q = 0$). In spherical geometry, points lie on the unit d-sphere $\mathcal{S}^d := \mathcal{S}_1^{d,0} = \{\mathbf{x} \in \mathbb{R}^{d+1} : \langle \mathbf{x}, \mathbf{x} \rangle = 1\}$. The geometry of the projective d-space $\mathcal{P}^d := \mathcal{S}^d / \pm 1$ is called elliptic geometry [24, 30]. Geodesic distances of $\mathcal{P}^d$ naturally account for the fact that they compare sets. Let $\overline{\mathsf{d}}_{\overline{\gamma}} : \mathcal{S}^d \times \mathcal{S}^d \to \mathbb{R}$ be the geodesic distance of $\mathcal{S}^d$ (i.e., spherical distance). The geodesic distance between $[\mathbf{x}] \in \mathcal{P}^d$ and $[\mathbf{y}] \in \mathcal{P}^d$ is $\mathsf{d}_\gamma([\mathbf{x}], [\mathbf{y}]) = \min_{\mathbf{a} \in [\mathbf{x}], \mathbf{b} \in [\mathbf{y}]} \overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{a}, \mathbf{b})$. We then have:

$$\mathsf{d}_\gamma([\mathbf{x}], [\mathbf{y}]) := \min\{\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, \mathbf{y}), \overline{\mathsf{d}}_{\overline{\gamma}}(-\mathbf{x}, \mathbf{y})\} = \cos^{-1}(|\langle \mathbf{x}, \mathbf{y} \rangle|) = \cos^{-1}(|\langle \mathbf{x}, \mathbf{y} \rangle_q|), \tag{4}$$

which is a distance metric. The fact that the spherical geometry is antipodally symmetric (i.e., every point can be inverted w.r.t. the origin) leads to a duplication of geometric information [24]. Identifying each pair of antipodal points to one point eliminates the antipodal duplication in spherical geometry.

The hyperboloid model of hyperbolic geometry is similar to the geometry of $\mathcal{P}_1^{0,q}$ ($p = 0$). The q-dimensional manifold $\mathcal{S}_1^{0,q} \subset \mathbb{R}^{1,q}$ contains two separate sheets (i.e., two connected components) and is anti-isometric to the hyperboloid of two sheets $\mathcal{Q}_1^{q,0}$. Pairs of antipodal points lying on different sheets of $\mathcal{S}_1^{0,q}$ are considered as a single point of $\mathcal{P}_1^{0,q}$. Let $\mathbf{x} \in \mathcal{S}_1^{0,q}$ and $\mathbf{z} \in \mathcal{S}_1^{0,q}$ be two points lying on the same sheet of $\mathcal{S}_1^{0,q}$, there exists no geodesic joining $\mathbf{x}$ and $-\mathbf{z}$. Their geodesic distance with respect to $\mathcal{S}_1^{0,q}$ can then be considered to be $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, -\mathbf{z}) = +\infty$, and we have:

$$\mathsf{d}_\gamma([\mathbf{x}], [\mathbf{z}]) := \min\{\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, \mathbf{z}), +\infty\} = \overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, \mathbf{z}) = \cosh^{-1}(\langle \mathbf{x}, \mathbf{z} \rangle_q) = \cosh^{-1}(|\langle \mathbf{x}, \mathbf{z} \rangle_q|), \tag{5}$$

which is similar to the hyperbolic distance metric of the hyperboloid model studied in [20].

Ultrahyperbolic geometry (or indefinite elliptic geometry). In this paper, we propose a parametric model that learns representations lying on the quotient manifold $\mathcal{P}_r^{p,q}$. When both p and q are positive, the metric tensor of $\mathcal{P}_r^{p,q}$ is nondegenerate (see page 343 of [16]) and indefinite. This means that the manifold is pseudo-Riemannian but not Riemannian due to the lack of positive definiteness of the metric tensor. $\mathcal{P}_r^{p,q}$ is also called an *indefinite elliptic space* [30] in the literature. We refer the reader to Chapters 11 and 12 of [30] or Chapter 7 of [21] for details. As an example, Fig. 1 illustrates the manifold $\mathcal{P}_1^{1,1}$. Our main motivation for considering $\mathcal{P}_r^{p,q}$ is that it is more flexible than hyperbolic and elliptic geometries since it contains hyperbolic and elliptic parts (i.e., time-like and space-like geodesics in Fig. 1). This flexibility allows us to better represent graphs that are not entirely trees or cycles, but that contain tree-like or cycle subgraphs. We experimentally verify our intuition.

3 Optimization on Ultrahyperbolic Quotient Manifolds

Our ultrahyperbolic representations lie on the quotient manifold $\mathcal{P}_r^{p,q}$. In this section, we provide differential geometry tools to optimize some differentiable function $f : \mathcal{P}_r^{p,q} \to \mathbb{R}$. To this end, we need the formulation of geodesics of $\mathcal{P}_r^{p,q}$. In Section 3.1, we explain how to formulate tangent vectors of $\mathcal{P}_r^{p,q}$ as a function of tangent vectors of $\mathcal{S}_r^{p,q}$ via the *horizontal lift* operator. This operator allows us to formulate geodesics of $\mathcal{P}_r^{p,q}$ as a function of geodesics of $\mathcal{S}_r^{p,q}$ in Section 3.2. In Section 3.3, we state the properties that the function f has to satisfy due to the quotient nature of $\mathcal{P}_r^{p,q}$. In Section 3.4, we illustrate how to optimize a standard neural network. Our deep GNN that maps activation representations in ultrahyperbolic space at each layer is introduced in Section 4.

3.1 Representing tangent vectors of $\mathcal{P}_r^{p,q}$ only by horizontal tangent vectors of $\mathcal{S}_r^{p,q}$

It can be difficult to work numerically with the tangent space $T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$ of $\mathcal{P}_r^{p,q}$ at $[\mathbf{x}]$ since $[\mathbf{x}] = \{\mathbf{x}, -\mathbf{x}\}$ is an equivalence class. We now present some differential geometry tools to define tangent vectors of $\mathcal{S}_r^{p,q}$ as a function of tangent vectors of $\mathcal{P}_r^{p,q}$, and vice versa. Their general definitions can be found in Chapter 7 of [21]. We also refer the reader to [4] for details on optimization on quotient manifolds. Our contribution in this subsection is that we give their formulation for $\mathcal{P}_r^{p,q}$. We first give the formulation of tangent spaces of $\mathcal{S}_r^{p,q}$ and then provide tools to identify tangent vectors of $\mathcal{P}_r^{p,q}$. These tools will be essential to construct geodesics of $\mathcal{P}_r^{p,q}$ and represent them via $\mathcal{S}_r^{p,q}$.

The tangent space $T_\mathbf{x}\mathcal{S}_r^{p,q}$ of $\mathcal{S}_r^{p,q}$ at $\mathbf{x}$ can be defined as: $T_\mathbf{x}\mathcal{S}_r^{p,q} := \{\boldsymbol{\xi} \in \mathbb{R}^{p+1,q} : \langle \boldsymbol{\xi}, \mathbf{x}\rangle_q = 0\}$.

The canonical map (or natural map [21]) $\pi : \mathcal{S}_r^{p,q} \to \mathcal{P}_r^{p,q}$ is defined as: $\forall \mathbf{x} \in \mathcal{S}_r^{p,q}, \pi(\mathbf{x}) := [\mathbf{x}] = \{\mathbf{x}, -\mathbf{x}\}$. Its differential at $\mathbf{x}$ is denoted by $\mathrm{d}\pi_\mathbf{x} : T_\mathbf{x}\mathcal{S}_r^{p,q} \to T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$.

The horizontal space $\mathcal{H}_\mathbf{x}$ and the vertical space $\mathcal{V}_\mathbf{x}$ at $\mathbf{x} \in \mathcal{S}_r^{p,q}$ are subspaces of $T_\mathbf{x}\mathcal{S}_r^{p,q}$ defined such that $T_\mathbf{x}\mathcal{S}_r^{p,q} = \mathcal{H}_\mathbf{x} \oplus \mathcal{V}_\mathbf{x}$ is a direct sum of linear spaces, and $\mathcal{V}_\mathbf{x}$ is the following kernel: $\mathcal{V}_\mathbf{x} := \ker(\mathrm{d}\pi_\mathbf{x})$. From Proposition 5.38 of [16], we find $\ker(\mathrm{d}\pi_\mathbf{x}) = T_\mathbf{x}([\mathbf{x}]) = \mathbf{0}$ because $[\mathbf{x}]$ is a discrete space so $[\mathbf{x}]$ and its tangent spaces are 0-dimensional. We then have $\mathcal{H}_\mathbf{x} = T_\mathbf{x}\mathcal{S}_r^{p,q}$. Elements of $\mathcal{H}_\mathbf{x}$ are called horizontal vectors, and all the tangent vectors of $\mathcal{S}_r^{p,q}$ are horizontal.

The horizontal lift (see §29 of [29]) at $\mathbf{x} \in \mathcal{S}_r^{p,q}$ of the tangent vector $\boldsymbol{\xi} \in T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$ is the unique horizontal vector denoted by $\overline{\boldsymbol{\xi}}_\mathbf{x} = \mathrm{lift}_\mathbf{x}(\boldsymbol{\xi}) \in \mathcal{H}_\mathbf{x}$ such that $\mathrm{d}\pi_\mathbf{x}(\overline{\boldsymbol{\xi}}_\mathbf{x}) = \boldsymbol{\xi}$. Since $\mathcal{H}_\mathbf{x} = T_\mathbf{x}\mathcal{S}_r^{p,q}$, the $\mathrm{lift}_\mathbf{x}$ operator is bijective so tangent vectors in $T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$ can be equivalently represented by horizontal vectors in $\mathcal{H}_\mathbf{x}$. During optimization, we will exploit this bijection and consider only some specific horizontal space to represent and update the weights of our neural network. The fact that $\mathcal{H}_\mathbf{x} = T_\mathbf{x}\mathcal{S}_r^{p,q}$ is convenient since it implies that any tangent vector in $T_\mathbf{x}\mathcal{S}_r^{p,q}$ can be represented in $T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$. We can then construct a geodesic of $\mathcal{P}_r^{p,q}$ from any geodesic of $\mathcal{S}_r^{p,q}$ as discussed below.

3.2 Geodesic of $\mathcal{P}_r^{p,q}$, exponential map and parallel transport

To optimize over $\mathcal{S}_r^{p,q}$ and $\mathcal{Q}_r^{q,p}$, Gao *et al.* [9] and Law & Stam [15] define tools such as the geodesic, parallel transport, exponential map, logarithm map and the *geodesic distance* $\overline{\mathrm{d}}_{\overline{\gamma}} : \mathcal{S}_r^{p,q} \times \mathcal{S}_r^{p,q} \to \mathbb{R}$ (see formulations in the supp. material). Our contribution in this subsection is that we extend all of the above differential geometry tools to $\mathcal{P}_r^{p,q}$. Their details can be found in the supp. material.

The geodesic $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}} : \mathbb{R} \to \mathcal{S}_r^{p,q}$ of $\mathcal{S}_r^{p,q}$ is the curve defined such that its initial point is $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(0) = \mathbf{x} \in \mathcal{S}_r^{p,q}$, its initial velocity is $\overline{\gamma}'_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(0) = \overline{\boldsymbol{\xi}}_\mathbf{x} \in T_\mathbf{x}\mathcal{S}_r^{p,q}$ and its acceleration is zero. When the initial conditions are clear from the context, we denote the geodesic by $\overline{\gamma}$ and ignore its indices. Since every geodesic $\overline{\gamma}$ of $\mathcal{S}_r^{p,q}$ satisfies $\forall t, \overline{\gamma}'(t) \in \mathcal{H}_{\overline{\gamma}(t)}$, it is called horizontal and $\gamma := \pi \circ \overline{\gamma} : \mathbb{R} \to \mathcal{P}_r^{p,q}$ is a geodesic of $\mathcal{P}_r^{p,q}$. By the chain rule, we have $\forall t, \gamma'(t) = \mathrm{d}\pi_{\overline{\gamma}(t)}(\overline{\gamma}'(t))$, which implies $\forall t, \mathrm{lift}_{\overline{\gamma}(t)}(\gamma'(t)) = \overline{\gamma}'(t)$. We then have $\forall t \in \mathbb{R}, \gamma_{[\mathbf{x}]\to\boldsymbol{\xi}}(t) = \{\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(t), \overline{\gamma}_{-\mathbf{x}\to\overline{\boldsymbol{\xi}}_{-\mathbf{x}}}(t)\}$, and we find $\overline{\boldsymbol{\xi}}_\mathbf{x} = -\overline{\boldsymbol{\xi}}_{-\mathbf{x}}$ to preserve the equivalence between antipodal points: $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(t) = -\overline{\gamma}_{-\mathbf{x}\to\overline{\boldsymbol{\xi}}_{-\mathbf{x}}}(t)$.

Exponential and logarithm map. The exponential map of $\mathcal{P}_r^{p,q}$ at $[\mathbf{x}]$ is the differentiable mapping $\exp_{[\mathbf{x}]} : T_{[\mathbf{x}]}\mathcal{P}_r^{p,q} \to \mathcal{P}_r^{p,q}$ defined such that $\exp_{[\mathbf{x}]}(\boldsymbol{\xi}) := \gamma_{[\mathbf{x}]\to\boldsymbol{\xi}}(1) = \{\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(1), \overline{\gamma}_{-\mathbf{x}\to\overline{\boldsymbol{\xi}}_{-\mathbf{x}}}(1)\}$.

We denote the exponential map of $\mathcal{S}_r^{p,q}$ at $\mathbf{x}$ by $\overline{\exp}_\mathbf{x} : T_\mathbf{x}\mathcal{S}_r^{p,q} \to \mathcal{S}_r^{p,q}$. It is defined as $\overline{\exp}_\mathbf{x}(\overline{\boldsymbol{\xi}}_\mathbf{x}) := \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_\mathbf{x}}(1)$, and we have $\exp_{[\mathbf{x}]}(\boldsymbol{\xi}) = [\overline{\exp}_\mathbf{x}(\overline{\boldsymbol{\xi}}_\mathbf{x})]$. In practice, we select some reference point $\mathbf{x}$ and only work with the exponential map $\overline{\exp}_\mathbf{x}$. The logarithm map is the inverse function of the exponential map (i.e., $\log_{[\mathbf{x}]} := \exp_{[\mathbf{x}]}^{-1}$). Their exact formulation can be found in the supp. material.

Parallel transport on $\mathcal{S}_r^{p,q}$. Given the minimizing (unbroken) geodesic $\overline{\gamma}$ (i.e., minimizing the arc length) from $\mathbf{x} = \overline{\gamma}(0)$ to $\mathbf{y} = \overline{\gamma}(1)$, the parallel transport $P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}} : T_\mathbf{x}\mathcal{S}_r^{p,q} \to T_\mathbf{y}\mathcal{S}_r^{p,q}$ is a linear isometry such that $\forall \overline{\boldsymbol{\xi}}_\mathbf{x}, \overline{\boldsymbol{\zeta}}_\mathbf{x},\ \langle \overline{\boldsymbol{\xi}}_\mathbf{x}, \overline{\boldsymbol{\zeta}}_\mathbf{x}\rangle_q = \langle P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_\mathbf{x}), P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\zeta}}_\mathbf{x})\rangle_q$ (see page 66 of [21]). The parallel transport along $\overline{\gamma}$ from $\mathbf{x}$ to $\mathbf{y}$ (where $\mathbf{x}$ and $\mathbf{y}$ satisfy $\langle \mathbf{x}, \mathbf{y}\rangle_q > -r^2$) is:

$$P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_\mathbf{x}) := \overline{\boldsymbol{\xi}}_\mathbf{x} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_\mathbf{x}\rangle_q}{\langle \mathbf{x}, \mathbf{y}\rangle_q + r^2}(\mathbf{y} + \mathbf{x}) \tag{6}$$

Minimizing geodesic of $\mathcal{P}_r^{p,q}$. Our parallel transport on $\mathcal{P}_r^{p,q}$ depends on a minimizing geodesic γ whose arc length (that we call *geodesic distance* d_γ) from $[\mathbf{x}] = \gamma(0)$ to $[\mathbf{y}] = \gamma(1)$ is:

$$\forall [\mathbf{x}] \in \mathcal{P}_r^{p,q}, [\mathbf{y}] \in \mathcal{P}_r^{p,q},\ \mathsf{d}_\gamma([\mathbf{x}],[\mathbf{y}]) = \begin{cases} r\cosh^{-1}(|\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}|) & \text{if } |\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}| \geq 1 \\ r\cos^{-1}\ (|\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}|) & \text{otherwise.} \end{cases} \tag{7}$$

and we have $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) < \overline{\mathsf{d}}_{\overline{\gamma}}(-\mathbf{x},\mathbf{y})$ iff $\langle \mathbf{x},\mathbf{y}\rangle_q > 0$. See details in the supp. material.

The parallel transport $P^{\gamma}_{[\mathbf{x}]\frown[\mathbf{y}]}$ on $\mathcal{P}_r^{p,q}$ can be horizontally lifted on $\mathcal{H}_\mathbf{y}$ as discussed above:

$$\forall \boldsymbol{\xi} \in T_{[\mathbf{x}]}\mathcal{P}_r^{p,q},\ \text{lift}_\mathbf{y}(P^{\gamma}_{[\mathbf{x}]\frown[\mathbf{y}]}(\boldsymbol{\xi})) = \begin{cases} P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_\mathbf{x}) & \text{if } \langle \mathbf{x},\mathbf{y}\rangle_q > 0 \\ P^{\overline{\gamma}}_{-\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_{-\mathbf{x}}) & \text{if } \langle \mathbf{x},\mathbf{y}\rangle_q < 0. \end{cases} \tag{8}$$

If $\langle \mathbf{x},\mathbf{y}\rangle_q = 0$, we have $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) = \overline{\mathsf{d}}_{\overline{\gamma}}(-\mathbf{x},\mathbf{y})$ and there exist two minimizing geodesics joining $[\mathbf{x}]$ and $[\mathbf{y}]$. In practice, we arbitrarily choose one of these two geodesics when $\langle \mathbf{x},\mathbf{y}\rangle_q = 0$.

3.3 Optimized function $f : \mathcal{P}_r^{p,q} \to \mathbb{R}$

Our goal is to minimize some differentiable function $f : \mathcal{P}_r^{p,q} \to \mathbb{R}$. We now describe the two properties that f has to satisfy. We first recall that every $[\mathbf{x}] \in \mathcal{P}_r^{p,q}$ is a set of equivalent elements that should preserve invariance. To simplify explanations, we consider the function $\overline{f} : \mathcal{S}_r^{p,q} \to \mathbb{R}$ defined such that $\overline{f} := f \circ \pi$. We then have $\forall \mathbf{x} \in \mathcal{S}_r^{p,q},\ \overline{f}(\mathbf{x}) = f([\mathbf{x}])$.

Property 1. Since $\mathbf{x}$ and $-\mathbf{x}$ are equivalent, the first property that f has to satisfy is $\overline{f}(\mathbf{x}) = \overline{f}(-\mathbf{x})$.

Property 2. Let $\nabla\overline{f}(\mathbf{x}) := (\partial\overline{f}(\mathbf{x})/\partial x_0, \dots, \partial\overline{f}(\mathbf{x})/\partial x_d)^\top$ be the Euclidean gradient of $\overline{f}$ at $\mathbf{x} = (x_0,\dots,x_d)^\top$. The pseudo-Riemannian gradient of $\overline{f}$ at $\mathbf{x} \in \mathcal{S}_r^{p,q}$ is $D\overline{f}(\mathbf{x}) := \Pi_\mathbf{x}(\mathbf{G}^{-1}\nabla\overline{f}(\mathbf{x})) = \Pi_\mathbf{x}(\mathbf{G}\nabla\overline{f}(\mathbf{x})) \in T_\mathbf{x}\mathcal{S}_r^{p,q}$ where $\Pi_\mathbf{x}(\mathbf{z}) := \mathbf{z} - \frac{\langle \mathbf{z},\mathbf{x}\rangle_q}{\langle \mathbf{x},\mathbf{x}\rangle_q}\mathbf{x}$ is the orthogonal projection of $\mathbf{z}$ onto $T_\mathbf{x}\mathcal{S}_r^{p,q}$.

Let $Df([\mathbf{x}]) \in T_{[\mathbf{x}]}\mathcal{P}_r^{p,q}$ be the pseudo-Riemannian gradient of f at $[\mathbf{x}] \in \mathcal{P}_r^{p,q}$. By applying the chain rule, the second property that f has to satisfy is $\text{lift}_\mathbf{x}(Df([\mathbf{x}])) = D\overline{f}(\mathbf{x}) = -D\overline{f}(-\mathbf{x})$.

3.4 Optimization of parametric models

We now explain how to minimize some function $f : \mathcal{P}_r^{p,q} \to \mathbb{R}$ that takes as input the ultrahyperbolic representation returned by some parametric model φ_θ (e.g., a neural network with parameters θ) that we want to learn. We exploit the fact that, due to the properties of the (affine) Levi-Civita connection [6, 17] of $\mathcal{P}_r^{p,q}$, the metric of the manifold $\mathcal{P}_r^{p,q}$ is preserved when we work with its tangent spaces via the exponential map (see page 61 of [21]).

Forward pass. Let us consider the positive pole $\mathbf{p} = (r, 0, \dots, 0)^\top \in \mathcal{S}_r^{p,q}$ defined such that only its first element $r > 0$ is nonzero. The horizontal space of $\mathbf{p}$ can be defined as the following vector space $\mathcal{H}_\mathbf{p} = T_\mathbf{p}\mathcal{S}_r^{p,q} = \{0\} \times \mathbb{R}^{p,q}$. The mapping $\varphi_\theta : \mathcal{X} \to \mathcal{H}_\mathbf{p}$ maps any input data $\boldsymbol{x} \in \mathcal{X}$ to $\mathcal{H}_\mathbf{p}$ and the resulting horizontal vector is mapped to $\mathcal{S}_r^{p,q}$ with the exponential map as follows: $\mathbf{x} := \overline{\exp}_\mathbf{p}(\varphi_\theta(\boldsymbol{x})) \in \mathcal{S}_r^{p,q}$. As mentioned above, working with the vector space $\mathcal{H}_\mathbf{p}$ greatly simplifies computations and preserves the metric thanks to the Levi-Civita connection of $\mathcal{P}_r^{p,q}$.

Note that for standard neural networks that map to $\mathbb{R}^d$, the tangent space is identified to the space itself by the natural isomorphism $T_\mathbf{x}\mathbb{R}^d \approx \mathbb{R}^d$ so the network weights also implicitly lie in the tangent space. Our approach extends Euclidean neural networks to $\mathcal{P}_r^{p,q}$.

Backward pass. We assume that the function $\overline{f} :\ \mathcal{S}_r^{p,q} \to \mathbb{R}$ satisfies the properties mentioned in Section 3.3. By exploiting Eq. (8), the horizontal lift of the parallel translate of the gradient $Df([\mathbf{x}])$ can be formulated as follows:

$$\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} := \text{lift}_\mathbf{p}\left(P^{\gamma}_{[\mathbf{x}]\frown[\mathbf{p}]}(Df([\mathbf{x}]))\right) = \begin{cases} P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{p}}(D\overline{f}(\mathbf{x})) & \text{if } \langle \mathbf{x},\mathbf{p}\rangle_q \geq 0 \\ P^{\overline{\gamma}}_{-\mathbf{x}\frown\mathbf{p}}(-D\overline{f}(\mathbf{x})) & \text{otherwise.} \end{cases} \tag{9}$$

Descent direction. When the metric tensor of the manifold is not positive definite, the manifold is not Riemannian and the negative of $\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$ is not a descent direction [9]. We show in the supp. material that the negative of $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} \in \mathcal{H}_\mathbf{p}$ is a descent direction that can be used to optimize the parameters of φ_θ with standard descent algorithms. We illustrate one such example in Section 5.1.

Complexity. Our optimizer exploits efficient closed-form expressions on $\mathcal{S}_r^{p,q}$ by considering $\mathbf{x}$ or its antipodal point $-\mathbf{x}$ depending on its "geodesic distance" with the positive pole $\mathbf{p}$. This geodesic distance depends only on the sign of $\langle \mathbf{x}, \mathbf{p} \rangle_q$, which is also the sign of the first element of $\mathbf{x} = (x_0, \dots, x_d)^\top$ (i.e., , $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, \mathbf{p}) < \overline{\mathsf{d}}_{\overline{\gamma}}(-\mathbf{x}, \mathbf{p}) \iff \langle \mathbf{x}, \mathbf{p} \rangle_q > 0 \iff x_0 > 0$). Our operators generalize tools used in hyperbolic space and are then as efficient as hyperbolic approaches.

4 Ultrahyperbolic Graph Convolutional Network (GCN)

We now extend the hyperbolic graph neural networks introduced in [18] to $\mathcal{P}_r^{p,q}$.

Graph Neural Networks. We first provide some background on Graph Neural Networks (GNNs) which can be interpreted as parametric models performing message passing between nodes of a graph. We recall the formulation of Graph Convolutional Networks (GCNs) [13] and rewrite them in our formalism with quotient manifolds. Let $G = (V, E)$ be a undirected graph containing $n = |V|$ nodes and $m = |E|$ edges. Its adjacency matrix is denoted by $\mathbf{A} \in \mathbb{R}^{n \times n}$. To account for self-loops, Liu *et al.* [18] consider the matrix $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2}(\mathbf{A} + \mathbf{I})\mathbf{D}^{-1/2}$ where $\mathbf{D}$ is the diagonal degree matrix defined such that $\mathbf{D}_{ii} = \sum_j (\mathbf{A}_{ij} + \mathbf{I}_{ij})$. The vector representation of node v at step k is denoted by $\mathbf{h}_v^k \in \mathbb{R}^d$, and $\mathbf{h}_v^0$ is given. $\mathbf{W}^k$ is a matrix whose elements are the trainable parameters of the k-th layer. The information in the Euclidean GCN propagates as: $\mathbf{h}_u^{k+1} = \sigma\left(\sum_{v \in \mathcal{I}(u)} \tilde{\mathbf{A}}_{uv} \mathbf{W}^k \mathbf{h}_v^k\right)$ where $\mathcal{I}(u)$ is the set of in-neighbors of $u \in V$ (i.e., u and v are joined by an edge) and σ is a nonlinear activation function such as the element-wise Rectified Linear Unit (ReLU) or its variants.

Ultrahyperbolic GNN. Let us now consider that $\forall v, k, \mathbf{h}_v^k \in \mathcal{P}_r^{p,q}$. Since $\mathcal{P}_r^{p,q}$ is not a vector space, the operation $\mathbf{W}^k \mathbf{h}_v^k$ is not defined, and the activation function σ has to be adapted. As in Section 3.4, we exploit properties of the Levi-Civita connection to work with the tangent spaces of $\mathcal{P}_r^{p,q}$ via the exponential map and its inverse (i.e., logarithm map). The propagation is then extended to $\mathcal{P}_r^{p,q}$ by:

$$\mathbf{h}_u^{k+1} := \sigma\left(\left[\overline{\exp}_{\mathbf{p}}\left(\sum_{v \in \mathcal{I}(u)} \tilde{\mathbf{A}}_{uv} \mathbf{W}^k \operatorname{lift}_{\mathbf{p}}\left(\log_{[\mathbf{p}]}(\mathbf{h}_v^k)\right)\right)\right]\right) \in \mathcal{P}_r^{p,q}, \tag{10}$$

where $\mathbf{p} = (r, 0, \dots, 0)^\top$ is the positive pole and we exploit the logarithm map to map points of $\mathcal{P}_r^{p,q}$ to a single tangent space. As explained in Section 3, in practice, we use the horizontal lift operator so that the exponential and logarithm maps only consider the horizontal space $\mathcal{H}_{\mathbf{p}}$ during optimization (see supp. material for details). The hyperbolic GNN [18] corresponds to the special case where $\mathcal{P}_r^{p,q} = \mathcal{P}_1^{0,q}$ (i.e., $p = 0$). We now give the formulation of the activation function σ.

Activation function via stereographic projection. For simplicity of exposition, we now consider that the radius of $\mathcal{S}_r^{p,q}$ is $r = 1$. To enforce nonlinearity between the different layers of the hyperbolic graph neural network, Liu et al. [18] formulate their activation function as the result of a steoreographic projection onto the negative pole $-\mathbf{p}$ from the hyperboloid model to the Poincaré ball, followed by a ReLU activation (in the Poincaré ball) and an inverse steoreographic projection from the Poincaré ball to the hyperboloid. We explain below how to generalize σ to pseudo-spheres.

Let us note $\varepsilon \in \{-1, 1\}$. The pole $\varepsilon\mathbf{p} = (\varepsilon, 0, \dots, 0)^\top$ is positive if $\varepsilon = 1$, and negative if $\varepsilon = -1$. Let us consider a point $\mathbf{x} = (x_0, x_1, \dots, x_d)^\top \in \mathcal{S}_1^{p,q}$ with $x_0 > 0$ (i.e., lying on the positive hemisphere). The stereographic projection of $\mathbf{x}$ onto $\varepsilon\mathbf{p}$ is $\mathbf{a} = \omega_\varepsilon(\mathbf{x}) := \frac{1}{1 - \varepsilon x_0}(x_1, x_2, \dots, x_d)^\top$. If $x_0 < 0$, we equivalently consider that $\mathbf{a} = \omega_\varepsilon(-\mathbf{x}) = -\omega_{-\varepsilon}(\mathbf{x})$ instead of $\omega_\varepsilon(\mathbf{x})$ due to the quotient nature of $\mathcal{P}_r^{p,q}$ and to account for the fact that $[\mathbf{x}]$ is projected onto the pole of different hemisphere if $\varepsilon = -1$, or same hemisphere if $\varepsilon = 1$. The inverse projection of $\mathbf{a} = (a_1, \dots, a_d)^\top \in \mathbb{R}^{p,q}$ is:

$$\omega_\varepsilon^{-1}(\mathbf{a}) := \frac{1}{1 + \langle \mathbf{a}, \mathbf{a} \rangle_q} \begin{pmatrix} \varepsilon(\langle \mathbf{a}, \mathbf{a} \rangle_q - 1) \\ 2\mathbf{a} \end{pmatrix} \in \mathcal{S}_1^{p,q} \quad \text{where } \langle \mathbf{a}, \mathbf{a} \rangle_q := \sum_{i=1}^{p} a_i^2 - \sum_{j=p+1}^{d} a_j^2. \tag{11}$$

We formulate $\sigma([\mathbf{x}]) := [\omega_\varepsilon^{-1}(\mathrm{ReLU}(\omega_\varepsilon(\mathbf{x})))]$ if $x_0 \geq 0$, and $\sigma([\mathbf{x}]) := [\omega_\varepsilon^{-1}(\mathrm{ReLU}(\omega_\varepsilon(-\mathbf{x})))]$ otherwise, where ReLU (or one of its variants such as LeakyRelu) is applied element-wise only on the q time dimensions of the input vector, which avoids having a zero denominator in Eq. (11). As in [18], we consider $\varepsilon = -1$. It is worth noting that Liu *et al.* [18] work with the upper sheet of the hyperboloid $\mathcal{Q}_1^{q,0}$ which is anti-isometric to $\mathcal{S}_1^{0,q}$. Their stereographic projection then contains only space dimensions. Their space dimensions correspond to our time dimensions due to anti-isometry.

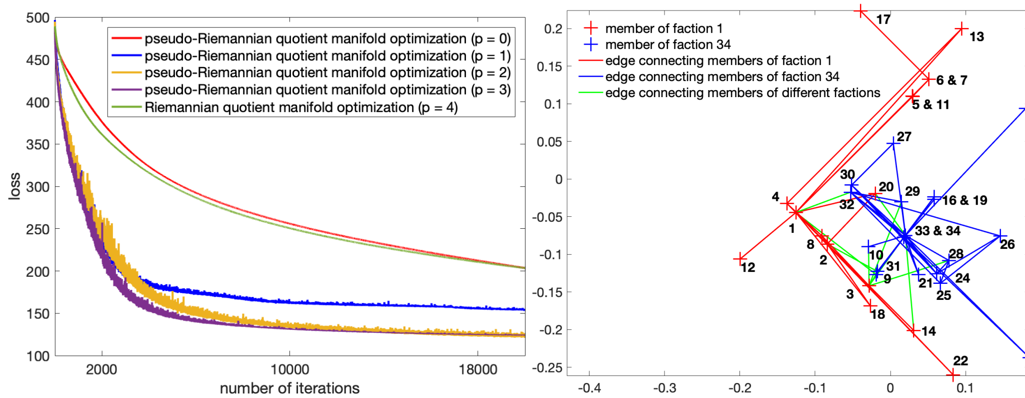

Figure 2: (left) Loss values of Eq. (12) as a function of the number of iterations for different values of p when $\mathcal{P}_r^{p,q}$ is 4-dimensional. (right) Stereographic projection onto $-\mathbf{p}$ of representations lying on $\mathcal{P}_1^{1,1}$ and learned from Zachary's karate club. Node colors define the faction joined by the members.

Table 1: Evaluation scores for the different learned representations (mean $\pm$ standard deviation)

Evaluation metric	$\mathbb{R}^4$ (Euclidean)	$\mathcal{P}_1^{0,4}$ (Hyperbolic)	$\mathcal{P}_1^{1,3}$	$\mathcal{P}_1^{2,2}$	$\mathcal{P}_1^{3,1}$	$\mathcal{P}_1^{4,0}$ (Elliptic)
Rank of first leader	4.6 ± 1.0	2.5 ± 0.7	$\mathbf{1.2 \pm 0.4}$	1.3 ± 0.7	$\mathbf{1.2 \pm 0.4}$	2.5 ± 0.8
Rank of second leader	6.9 ± 0.7	3.8 ± 1.0	$\mathbf{2.7 \pm 0.7}$	3.1 ± 1.0	4.4 ± 3.0	3.6 ± 0.7
top 5 Spearman's ρ	0.06 ± 0.45	0.36 ± 0.22	0.62 ± 0.23	0.61 ± 0.28	$\mathbf{0.63 \pm 0.35}$	0.46 ± 0.29
top 10 Spearman's ρ	0.04 ± 0.19	0.38 ± 0.18	$\mathbf{0.73 \pm 0.12}$	0.72 ± 0.07	0.63 ± 0.16	0.38 ± 0.26
Training time (seconds)	$\mathbf{340 \pm 4}$	424 ± 1	429 ± 1	430 ± 2	429 ± 1	402 ± 1

5 Experiments

We now evaluate our approach on different classification tasks on graphs. We first show that our optimization framework introduced in Section 3.4 learns meaningful representations on a toy hierarchical graph with cycles. We then apply our framework in standard classification tasks.

5.1 Last layer optimization on a toy dataset

We evaluate our optimization framework by training a multi-layer perceptron (MLP) $\varphi_\theta : \mathcal{X} \to \mathcal{H}_\mathbf{p}$ whose set of parameters is called θ. As in [15], we test our approach on Zachary's karate club dataset [33]. However, instead of learning embeddings, we train a parametric model.

Zachary's dataset is a social network graph that represents a karate club split in two factions due to a conflict between two leaders (the instructor and the administrator). It is an undirected graph $G = (V, E)$ which has node-set $V = \{v_i\}_{i=1}^n$ and edge-set $E = \{e_k\}_{k=1}^m$ where $n = 34$ and $m = 78$. Each node v_i represents a karate member and an edge joins two nodes if the two members are friends. The two leaders are v_1 and v_{34}. We consider that each node v_i is represented as a distinct n-dimensional one-hot vector $\boldsymbol{x}_i \in \mathcal{X}$.

Problem. Following [15], our goal is to learn representations of nodes such that pairs of nodes joined by an edge (i.e., in E) have smaller distance than pairs of nodes that are not joined by an edge (i.e., not in E). Our problem is then to find the set of parameters θ that minimizes the problem:

$$\min_\theta \sum_{(v_i,v_j)\in E} -\log \frac{e^{-\mathsf{d}(\varrho_\theta(\boldsymbol{x}_i),\varrho_\theta(\boldsymbol{x}_j))/\tau}}{\sum_{(v_a,v_b)\in\, \mathcal{W}_{ij}} e^{-\mathsf{d}(\varrho_\theta(\boldsymbol{x}_a),\varrho_\theta(\boldsymbol{x}_b))/\tau}} \text{ where } \varrho_\theta(\boldsymbol{x}_i) := [\overline{\exp}_\mathbf{p}(\varphi_\theta(\boldsymbol{x}_i))] \quad (12)$$

and where $\mathcal{W}_{ij} := \{(v_i, v_j)\} \cup \{(v_a, v_b) \notin E\}$, $\tau = 10^{-2}$ is a fixed temperature value, and d denotes the geodesic distance of the manifold (e.g., Eq. (7) for $\mathcal{P}_r^{p,q}$). The geodesic distance satisfies the two properties defined in Section 3.3 with respect to each input and can then be used for optimization.

Model. Our MLP $\varphi_\theta : \mathcal{X} \to \mathcal{H}_\mathbf{p}$ contains three hidden layers of 10^4 hidden units each, with standard ReLU as nonlinear activation function. In this toy experiment, our MLP is standard, with the only exception that its last layer maps to the horizontal space $\mathcal{H}_{\mathbf{p}} = \{0\} \times \mathbb{R}^{p,q}$ of the positive pole $\mathbf{p}$. The output representation is then mapped with the exponential map as explained in Section 3.4.

Optimizer. We use the optimizer introduced in Section 3.4 to update θ. By using the descent direction $-\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}_i],\mathbf{p}}$ for each sample $[\mathbf{x}_i] = \varrho_\theta(\boldsymbol{x}_i)$, all the parameters of our standard MLP lie in some space equipped with a positive definite metric tensor. Standard backpropagation is then used to optimize the parameters. As an illustration, we consider the 4-dimensional manifold $\mathcal{P}_1^{p,q}$ (i.e., $p + q = 4$) and show in Fig. 2 (left) the loss values of Eq. (12) as a function of the number of iterations for different values of $p \in \{0, \dots, 4\}$. The figure shows that the optimization framework in Section 3.4 decreases the loss value. Moreover, it is worth noting that the algorithm does not converge if $-\boldsymbol{\lambda}_{[\mathbf{x}_i],\mathbf{p}}$ is used as a search direction (instead of $-\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}_i],\mathbf{p}}$) when the metric tensor is not positive definite since $-\boldsymbol{\lambda}_{[\mathbf{x}_i],\mathbf{p}}$ is not a descent direction [9]. More details can be found in the supp. material.

Hierarchy extraction. We now evaluate the quality of the learned representations in the task of predicting the high-level nodes of the graph. Our evaluation protocol is similar to [15], the only difference is that we train a neural network. We run 10 random initializations for each considered 4-dimensional manifold and report in Table 1 the mean and standard deviation of the different evaluation metrics.

As in [15], following the idea that hyperbolic distances grow exponentially, we take the sum of distances $\delta_i = \sum_{j=1}^{n} \mathsf{d}([\mathbf{x}_i], [\mathbf{x}_j])$ of a node v_i with all the other nodes as an indicator of importance. We sort the different $\delta_1, \dots, \delta_n$ in ascending order and report the rank of the two leaders (instructor and administrator, in no particular order) in the first two rows of Table 1. The leaders tend to have smaller δ_i score than low-level nodes with ultrahyperbolic distances, which means that high-level nodes tend to be closer to the rest of the nodes in ultrahyperbolic space.

We also measure the Spearman's rank correlation coefficient [28] between the 5 (or 10) most important nodes in the hierarchy and their corresponding δ_i score. Once again, the order of the δ_i scores is more correlated with the hierarchy level in ultrahyperbolic space. Our experimental results are comparable with [15] although our nodes are represented on a quotient manifold and we learn a parametric model. Fig. 2 (right) illustrates our learned representations when the manifold is $\mathcal{P}_1^{1,1}$.

Products of Riemannian space forms. In Table 1, we compare the performance of models mapping representations to pseudo-Riemannian space forms (i.e., manifolds of constant curvature [21, 30]). Nonetheless, it was already noticed in the machine learning literature that products of Riemannian space forms (called *mixed-curvature representations*) could outperform Riemannian space forms when the structure of the dataset is not tree-like [3, 11]. It is worth noting that products of space forms are in general not space forms (except if they are all flat). For this reason, we do not compare them to our manifold in the main article as we could similarly consider products of pseudo-spheres $\mathcal{P}_{r_1}^{p_1,q_1} \times \mathcal{P}_{r_2}^{p_2,q_2}$ or even $\mathcal{P}_{r_1}^{p_1,q_1} \times \mathbb{R}^{p_2,q_2}$ for evaluation.

Nonetheless, since our space form $\mathcal{P}_r^{p,q}$ contains hyperbolic and elliptic parts, we provide a detailed comparison with products of hyperbolic and spherical spaces in the supp. material. Such product manifolds perform better than hyperbolic and spherical spaces but slightly worse than the pseudo-Riemannian space form $\mathcal{P}_r^{p,q}$.

Training times. We report in Table 1 the training times of our Pytorch [22] implementation to train 25,000 iterations on a machine equipped with a 6-core Intel i7-7800X CPU and NVIDIA GeForce RTX 3090 GPU. All the representations lying on a non-flat manifold have comparable training times. Nonetheless, they are 25% slower than the Euclidean approach because they compute the pseudo-Riemannian gradient (which requires an orthogonal projection) and parallel transport.

5.2 Classification with ultrahyperbolic graph convolutional networks

The previous subsection analyzed our framework. We now evaluate it in standard classification tasks.

Node classification. We now evaluate the generalization performance of our GCN in the semi-supervised node classification task on three citation network datasets: Citeseer, Cora and Pubmed [26]. They contain sparse bag-of-words feature vectors for each document and a list of citation links between documents. Each document is a node and has a class label. Each citation link is an undirected edge. Dataset statistics are reported in Table 2. During training, all the nodes and edges are preserved, but only 20 nodes per class are labeled, and 500 nodes are used for validation in total, the rest for test. We follow the experimental protocol of Appendix A of [18] and learn a GCN with 2 hidden layers.

Table 2: Statistics of the citation network datasets.

Name	# Nodes	# Edges	# Classes	# Features	# training nodes per category
Citeseer	3,327	4,732	6	3,703	20
Cora	2,708	5,429	7	1,433	20
Pubmed	19,717	44,338	3	500	20

Table 3: Test node classification accuracy with 4-dimensional manifolds

Dataset	$\mathbb{R}^4$ (standard GCN)	$\mathcal{P}_1^{0,4}$ (Hyperbolic)	$\mathcal{P}_1^{1,3}$	$\mathcal{P}_1^{2,2}$	$\mathcal{P}_1^{3,1}$	$\mathcal{P}_1^{4,0}$ (Elliptic)
Citeseer	44.5 ± 5.9	46.7 ± 1.8	**51.8 ± 2.6**	50.3 ± 2.1	51.4 ± 3.2	47.2 ± 2.6
Cora	53.5 ± 4.3	56.2 ± 3.1	63.2 ± 3.3	63.9 ± 3.1	**64.7 ± 5.3**	61.4 ± 1.5
Pubmed	66.9 ± 2.3	71.5 ± 2.9	**73.1 ± 0.6**	72.8 ± 2.7	71.2 ± 2.7	71.0 ± 2.7

Table 4: Statistics of the graph datasets used for the classification task

Name	# graphs	# classes	Avg. # nodes	Avg. # edges	Type of dataset
Collab	5,000	3	74.49	2457.78	Scientific collaboration dataset [32]
D&D	1,178	2	284.32	715.66	Protein dataset [25]
Enzymes	600	6	32.63	62.14	Protein dataset [25]
Proteins	1,113	2	39.06	72.82	Protein dataset [25]
Reddit-multi-12K	11,929	11	391.41	456.89	Social network dataset [32]

Table 5: Graph classification accuracy in percents. d is the dimensionality of the manifold.

Method	Collab ($d = 64$)	D&D ($d = 88$)	Enzymes ($d = 256$)	Proteins ($d = 100$)	Reddit ($d = 100$)
Euclidean (standard GCN)	81.88 ± 1.76	76.93 ± 7.21	43.83 ± 10.3	75.46 ± 3.88	45.65 ± 1.76
Poincaré (hyperbolic)	80.92 ± 1.99	75.89 ± 8.53	44.15 ± 8.43	73.64 ± 4.64	45.84 ± 1.42
Lorentz (hyperbolic)	81.32 ± 1.21	77.10 ± 6.65	44.83 ± 8.14	74.16 ± 3.25	45.39 ± 1.53
Ultrahyperbolic	**82.26 ± 1.23**	**81.97 ± 3.41**	**50.50 ± 6.71**	**76.56 ± 2.09**	**47.08 ± 1.26**

When the dimensionality of each layer is $d = 600$, all the Euclidean (i.e., standard), Hyperbolic and Ultrahyperbolic GCNs reach the same test accuracy because the model is overparameterized and quickly attains 100% accuracy on the training set. See details and scores in the supp. material.

Due to the problem mentioned above, we trained GCNs whose dimensionality of each layer is $d = 4$ with 100 random initializations. The results reported in Table 3 show the superiority of ultrahyperbolic representations in low-dimensional space for node classification of hierarchical graphs with cycles. We also report results for $d = 10$ in the supp. material. The conclusion is similar.

Graph classification. We also evaluate our approach on commonly used graph kernel benchmark datasets [12] whose statistics are reported in Table 4. The evaluation is done via 10-fold cross validation. We use the same protocol evaluation and splits as in Appendix E of [18] and evaluate our approach in the same settings including same number of GNN layers, optimizers, learning rate, and manifold dimensionality d reported in Table 5. The only difference is that the data is represented on $\mathcal{P}_r^{p,q}$ with $p = 1$ in our case. The comparative performances are reported in Table 5 and show that ultrahyperbolic representations significantly improve performance on the D&D and Enzymes datasets, which are protein datasets from [25]. The gain is less significant on the other datasets but our approach is still competitive. It seems that the advantage of our approach over hyperbolic approaches is more visible for protein structures than for social networks, at least in high-dimensional space. More details can be found in the supp. material.

6 Conclusion, Limitations and Potential Societal Impacts

We have introduced neural networks that map data to a (quotient) pseudo-Riemannian manifold of constant nonzero curvature. Our considered geometry generalizes both hyperbolic and elliptic geometries. It is the first neural network that maps data to a non-Riemannian manifold to the best of our knowledge. Our framework is general and can be applied to many parametric models and tasks. We demonstrate this via graph convolutional networks and show improved performance compared to Euclidean and hyperbolic approaches to represent hierarchical graphs in different tasks.

Concurrently with this work, Xiong *et al.* [31] proposed an extension of graph convolutional networks to the pseudo-hyperboloid $\mathcal{Q}_r^{q,p}$ which is a pseudo-Riemannian manifold of constant nonzero curvature anti-isometric to the pseudo-sphere $\mathcal{S}_r^{p,q}$. One main difference is that, since there exist pairs of points of $\mathcal{Q}_r^{q,p}$ that cannot be joined by an unbroken geodesic, the optimization framework in [31] does not exploit the intrinsic geometry of the manifold via its Levi-Civita connection. On the other hand, our approach uses pseudo-Riemannian optimization tools that are intrinsic to $\mathcal{P}_r^{p,q}$. The ablation study in [31] also suggests that graphs with more hierarchical structure are better represented when the manifold becomes more hyperbolic, and graphs with cyclic relationships are better represented when the manifold becomes more spherical.

Limitations. Our main contribution is a solid optimization framework that is well defined thanks to the use of standard differential geometry tools (e.g., canonical map and horizontal bundle) that we formulate for the quotient manifold $\mathcal{P}_r^{p,q}$. It only requires the properties of the optimized function in Section 3.3 to be satisfied. This is for instance the case if points of $\mathcal{P}_r^{p,q}$ are compared with the geodesic distance in Eq. (7). We applied our framework on nine different datasets with (at least 10) different runs to validate our results. Our work lacks a theoretical analysis similar to Gromov's work [10] in the case of graphs without cycles. However, the optimal geometry for graphs with cycles is still an open problem, and hyperbolic geometry is used heuristically in this case. Our motivation is that ultrahyperbolic manifolds are more general than hyperbolic and elliptic manifolds, they can then combine the strengths of the two induced geometries. We experimentally validate our assumption in different tasks and leave the theoretical analysis for future work.

Potential societal impacts. Our contributions are mainly methodological although we apply our approach to hierarchical graphs that could represent social networks. Improving accuracy on these datasets might facilitate the task of discovering leaders in social networks, which could have negative impact if not monitored. Nonetheless, we also show improvement on protein structures, this could have positive impacts on society and healthcare. We did not exploit any personally identifiable information. We used datasets that have been publicly available to the machine learning community for years. Our method to handle and process the data is standard in the graph community.

Acknowledgments and Funding Transparency Statement. I thank James Lucas, Rafid Mahmood, Haggai Maron and the anonymous reviewers for helpful feedback on early versions of this manuscript. This project was entirely funded by NVIDIA corporation while I was working from home during the COVID-19 pandemic. I am also grateful to the NeurIPS 2021 program committee for giving me free registration to the conference thanks to an Outstanding Reviewer Award.

References

[1] Pierre-Antoine Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.

[2] Henri Anciaux. *Minimal submanifolds in pseudo-Riemannian geometry*. World Scientific, 2011.

[3] Gregor Bachmann, Gary Bécigneul, and Octavian Ganea. Constant curvature graph convolutional networks. In *International Conference on Machine Learning*, pages 486–496. PMLR, 2020.

[4] Nicolas Boumal. An introduction to optimization on smooth manifolds. Available online, Nov 2020.

[5] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. In *Advances in neural information processing systems*, pages 4868–4879, 2019.

[6] EB Christoffel. Ueber die transformation der homogenen differentialausdrücke zweiten grades. *Journal für die reine und angewandte Mathematik*, 70:46–70, 1869.

[7] Albert Einstein. Zur elektrodynamik bewegter körper. *Annalen der physik*, 322(10):891–921, 1905.

[8] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *Advances in neural information processing systems*, pages 5345–5355, 2018.

[9] Tingran Gao, Lek-Heng Lim, and Ke Ye. Semi-riemannian manifold optimization. *arXiv preprint arXiv:1812.07643*, 2018.

[10] Mikhael Gromov. Hyperbolic groups. In *Essays in group theory*, pages 75–263. Springer, 1987.

[11] Albert Gu, Frederic Sala, Beliz Gunel, and Christopher Ré. Learning mixed-curvature representations in product spaces. In *International Conference on Learning Representations*, 2018.

[12] Kristian Kersting, Nils M. Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. Benchmark data sets for graph kernels, 2016.

[13] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2017.

[14] Marc T. Law, Renjie Liao, Jake Snell, and Richard S. Zemel. Lorentzian distance learning for hyperbolic representations. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3672–3681, 2019.

[15] Marc T. Law and Jos Stam. Ultrahyperbolic representation learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 1668–1678, 2020.

[16] John M. Lee. *Introduction to Smooth Manifolds*, volume 218. Springer Science & Business Media, 2013.

[17] Tullio Levi-Civita. Nozioni di parallelismo in una varieta qualunque e consequente specificazione geometrica della curvature riemannia. *Rendiconti del Circolo Matematico di Palermo*, 42:173–205, 1917.

[18] Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. In *Advances in Neural Information Processing Systems 32*, pages 8228–8239, 2019.

[19] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*, pages 6338–6347, 2017.

[20] Maximillian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International Conference on Machine Learning*, pages 3779–3788, 2018.

[21] Barrett O'neill. *Semi-Riemannian geometry with applications to relativity*, volume 103. Academic press, 1983.

[22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[23] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. Hyperbolic deep neural networks: A survey. *arXiv preprint arXiv:2101.04562*.

[24] John Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149. Springer Science & Business Media, 2006.

[25] Ida Schomburg, Antje Chang, Christian Ebeling, Marion Gremse, Christian Heldt, Gregor Huhn, and Dietmar Schomburg. Brenda, the enzyme database: updates and major new developments. *Nucleic acids research*, 32(suppl_1):D431–D433, 2004.

[26] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.

[27] Ryohei Shimizu, Yusuke Mukuta, and Tatsuya Harada. Hyperbolic neural networks++. In *International Conference on Learning Representations*, 2021.

[28] Charles Spearman. The proof and measurement of association between two things. *The American journal of psychology*, 1904.

[29] Loring W Tu. *Differential geometry: connections, curvature, and characteristic classes*, volume 275. Springer, 2017.

[30] Joseph Albert Wolf. *Spaces of constant curvature*, volume 372. AMS Chelsea Publishing: An Imprint of the American Mathematical Society, Sixth Edition, 2011.

[31] Bo Xiong, Shichao Zhu, Nico Potyka, Shirui Pan, Chuan Zhou, and Steffen Staab. Semi-riemannian graph convolutional networks. *arXiv preprint arXiv:2106.03134*, 2021.

[32] Pinar Yanardag and SVN Vishwanathan. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1365–1374. ACM, 2015.

[33] Wayne W. Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, 33(4):452–473, 1977.

A Supplementary Material

The supplementary material is structured as follows:

• In Section B, we give the formulations of the differential geometry tools to work on the pseudo-sphere $\mathcal{S}_r^{p,q}$ (Section B.1), the indefinite elliptic space $\mathcal{P}_r^{p,q}$ (Section B.2) and the pseudo-hyperboloid $\mathcal{Q}_r^{q,p}$ (Section B.3). The tools include the formulation of a geodesic, exponential map, logarithm map and geodesic *distance*. In Section B.4, we explain the anti-isometry between the pseudo-sphere and the pseudo-hyperboloid. In Section B.5, we give more details about Figure 1. In Section B.6, we explain how the ultrahyperbolic manifold $\mathcal{P}_r^{p,q}$ contains hyperbolic and spherical/elliptic parts.

• In Section C, we explain how we optimize our neural networks. In particular, the pseudo-Riemannian gradient is not always a descent direction so we exploit results in [9] to find a descent direction in an efficient way.

• In Section D, we provide experimental details and additional results.

B Differential Geometry Tools

We provide here the necessary differential geometry tools to work on the pseudo-sphere $\mathcal{S}_r^{p,q}$ and the quotient manifold $\mathcal{P}_r^{p,q}$. Most of them are explained in [15] for the case of the pseudo-hyperboloid that is anti-isometric to the pseudo-sphere (see Section B.4 for details). We recall that the radius r of the pseudo-sphere is positive, and we consider that $r = 1$ in our experiments.

B.1 Pseudo-sphere $\mathcal{S}_r^{p,q}$

We give here the differential geometry tools specific to the pseudo-sphere which is defined as the following set: $\mathcal{S}_r^{p,q} := \left\{\mathbf{x} \in \mathbb{R}^{p+1,q} : \langle \mathbf{x}, \mathbf{x} \rangle_q = r^2\right\}$.

B.1.1 Geodesic, exponential map and distance

Geodesic. The geodesic $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}} : \mathbb{R} \to \mathcal{S}_r^{p,q}$ satisfying $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(0) = \mathbf{x}$ and $\overline{\gamma}'_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(0) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ is formulated for all $t \in \mathbb{R}$:

$$\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(t) = \begin{cases} \cos\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}} \sin\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q > 0 \\ \mathbf{x} + t\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q = 0 \\ \cosh\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}} \sinh\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q < 0 \end{cases} \tag{13}$$

The nonconstant geodesic $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}$ (i.e., $\overline{\boldsymbol{\xi}}_{\mathbf{x}} \neq \mathbf{0}$) is called:

• **space-like** if $\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q > 0$.

• **null** if $\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q = 0$.

• **time-like** if $\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q < 0$.

Exponential map. The exponential map $\overline{\exp}_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{S}_r^{p,q} \to \mathcal{S}_r^{p,q}$ is defined such that $\forall \overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}, \overline{\exp}_{\mathbf{x}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) = \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(1)$. We then have:

$$\overline{\exp}_{\mathbf{x}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) = \begin{cases} \cos\left(\frac{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}} \sin\left(\frac{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q > 0 \\ \mathbf{x} + \overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q = 0 \\ \cosh\left(\frac{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}} \sinh\left(\frac{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle\overline{\boldsymbol{\xi}}_{\mathbf{x}},\overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q < 0 \end{cases} \tag{14}$$

Logarithm map. The logarithm map $\overline{\log}_{\mathbf{x}}$ is defined as the inverse of the exponential map $\overline{\exp}_{\mathbf{x}}$ on a normal neighborhood of $\mathbf{x} \in \mathcal{S}_r^{p,q}$ denoted by $\mathcal{U}_{\mathbf{x}} = \{\mathbf{y} \in \mathcal{S}_r^{p,q} : \frac{\langle\mathbf{x},\mathbf{y}\rangle_q}{r^2} > -1\}$. It is then formulated:

$$\forall \mathbf{y} \in \mathcal{U}_{\mathbf{x}},\ \overline{\log}_{\mathbf{x}}(\mathbf{y}) = \begin{cases} \frac{\cos^{-1}(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2})}{\sqrt{1-(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2})^2}} \left(\mathbf{y} - \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}\mathbf{x}\right) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2} \in (-1,1) \\ \mathbf{y}-\mathbf{x} & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2} = 1 \\ \frac{\cosh^{-1}(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2})}{\sqrt{(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2})^2-1}} \left(\mathbf{y} - \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}\mathbf{x}\right) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2} > 1 \end{cases} \tag{15}$$

Geodesic "distance". As explained in [15] and Chapter 5 of [21], when the logarithm map $\overline{\log}$ exists for some pseudo-Riemannian manifold $\mathcal{M}$, the arc length of the tangent vector joining $\mathbf{x} \in \mathcal{M}$ and $\mathbf{y} \in \mathcal{M}$ corresponds to the radius function: $\sqrt{|g_{\mathbf{x}}(\overline{\log}_{\mathbf{x}}(\mathbf{y}), \overline{\log}_{\mathbf{x}}(\mathbf{y}))|}$ where $g_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{M} \times T_{\mathbf{x}}\mathcal{M} \to \mathbb{R}$ is the metric tensor at $\mathbf{x}$ and $\overline{\log}_{\mathbf{x}}$ is the logarithm map. In the case of the pseudo-sphere, we have $g_{\mathbf{x}}(\cdot,\cdot) = \langle \cdot,\cdot\rangle_q$. The geodesic distance $\overline{\mathsf{d}}_{\overline{\gamma}} : \mathcal{S}_r^{p,q} \times \mathcal{S}_r^{p,q} \to \mathbb{R}$ is then:

$$\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) = \sqrt{|\langle \overline{\log}_{\mathbf{x}}(\mathbf{y}), \overline{\log}_{\mathbf{x}}(\mathbf{y})\rangle_q|} = \begin{cases} r\cosh^{-1}(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2} \geq 1 \\ r\cos^{-1}(\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2} \in (-1,1) \end{cases} \tag{16}$$

$\overline{\mathsf{d}}_{\overline{\gamma}}$ is not a "distance metric" but a symmetric premetric: it satisfies (i) $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) = \overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{y},\mathbf{x}) \geq 0$ and (ii) $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{x}) = 0$.

In [21], the "minimizing geodesic" is defined by its arc length and then also corresponds to our geodesic distance.

B.1.2 Parallel transport on $\mathcal{S}_r^{p,q}$

The parallel transport formula is given in Eq. (6) of the main paper. For completeness, we write it here again. We also provide the proof that is inspired by [9] wherein the parallel transport on $\mathcal{S}_1^{p,q}$ along any geodesic is provided. We assume that $\mathbf{x}$ and $\mathbf{y}$ can be joined by an unbroken geodesic, the minimizing geodesic can then be formulated as a function of the logarithm map.

Given the minimizing geodesic $\overline{\gamma}$ connecting $\mathbf{x}$ to $\mathbf{y}$, the parallel transport $P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}} : T_{\mathbf{x}}\mathcal{S}_r^{p,q} \to T_{\mathbf{y}}\mathcal{S}_r^{p,q}$ is a linear isometry such that $\forall \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}},\ \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q = \langle P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}), P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\zeta}}_{\mathbf{x}})\rangle_q$. The parallel transport along $\overline{\gamma}$ from $\mathbf{x} = \overline{\gamma}(0)$ to $\mathbf{y} = \overline{\gamma}(1)$ (where $\mathbf{x}$ and $\mathbf{y}$ satisfy $\langle \mathbf{x},\mathbf{y}\rangle_q > -r^2$) is:

$$P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) := \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q}{\langle \mathbf{x},\mathbf{y}\rangle_q + r^2}(\mathbf{y}+\mathbf{x}) \tag{17}$$

Proof. To prove the correctness of the above formula, we follow the general properties of parallel transport mentioned in [9]. We briefly recall them here. We refer the reader to [9] for details.

We denote the *semi-normal space* of $\mathcal{S}_r^{p,q}$ in $\mathbb{R}^{p+1,q}$ at $\mathbf{x}$ by $\mathrm{SN}_{\mathbf{x}}(\mathcal{S}_r^{p,q}, \mathbb{R}^{p+1,q})$. It is defined as:

$$\mathrm{SN}_{\mathbf{x}}(\mathcal{S}_r^{p,q}, \mathbb{R}^{p+1,q}) := \{\mathbf{y} \in \mathbb{R}^{p+1,q} : \forall \overline{\boldsymbol{\zeta}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}, \langle \mathbf{y}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q = 0\} = \{\lambda\mathbf{x} : \lambda \in \mathbb{R}\} \tag{18}$$

A parallel translation of $\overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ along some geodesic $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}} : \mathbb{R} \to \mathcal{S}_r^{p,q}$ is a vector field. For the purpose of notation, we write $\overline{\gamma}$ instead $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}$ when the indices are not necessary and this vector field satisfies $\overline{\boldsymbol{\xi}}_{\mathbf{x}}(0) = \overline{\boldsymbol{\xi}}_{\mathbf{x}}$ and $\forall t \in \mathbb{R}, \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t) := \frac{\mathrm{D}}{dt}\left(\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t)\right) \in \mathrm{SN}_{\overline{\gamma}(t)}(\mathcal{S}_r^{p,q}, \mathbb{R}^{p+1,q})$ where $\frac{\mathrm{D}}{dt}\left(\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t)\right)$ is the covariant derivative of $\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t)$ along $\overline{\gamma}(t)$ in the ambient space $\mathbb{R}^{p+1,q}$.

By definition of the parallel transport, we have $\forall t, \overline{\boldsymbol{\xi}}_{\mathbf{x}}(t) \in T_{\overline{\gamma}(t)}\mathcal{S}_r^{p,q}$, which implies:

$$\forall t, \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}(t), \overline{\gamma}(t)\rangle_q = 0 \implies \langle \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t), \overline{\gamma}(t)\rangle_q = -\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}(t), \overline{\gamma}'(t)\rangle_q \quad \text{(obtained by differentiating)} \tag{19}$$

By definition, we have $\forall t, \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t) \in \mathrm{SN}_{\overline{\gamma}(t)}(\mathcal{S}_r^{p,q}, \mathbb{R}^{p+1,q})$ and $\forall t, \frac{1}{r^2}\langle \overline{\gamma}(t), \overline{\gamma}(t)\rangle_q = 1$, which implies $\forall t, \langle \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t), \overline{\gamma}(t)\rangle_q = \frac{1}{r^2}\langle \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t), \overline{\gamma}(t)\rangle_q \langle \overline{\gamma}(t), \overline{\gamma}(t)\rangle_q$ and we have:

$$\forall t, \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t) = \frac{1}{r^2}\langle \dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t), \overline{\gamma}(t)\rangle_q \overline{\gamma}(t) = -\frac{1}{r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}(t), \overline{\gamma}'(t)\rangle_q \overline{\gamma}(t) \tag{20}$$

Since parallel translation preserves the metric, we have $\forall t, \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}(t), \overline{\gamma}'(t)\rangle_q = \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}(0), \overline{\gamma}'(0)\rangle_q = \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q$. By using the initial condition $\overline{\boldsymbol{\xi}}_{\mathbf{x}}(0) = \overline{\boldsymbol{\xi}}_{\mathbf{x}}$ and integrating $\dot{\overline{\boldsymbol{\xi}}}_{\mathbf{x}}(t) = -\frac{1}{r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q \overline{\gamma}(t)$, the parallel transport of $\overline{\boldsymbol{\xi}}_{\mathbf{x}}$ along $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(t)$ is:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t) := \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{1}{r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q \int_0^t \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(\tau) d\tau \tag{21}$$

We have three cases to consider:

(1) If $\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q = 0$, we find (see Eq. (13)):

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{1}{r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q (t\mathbf{x} + \frac{1}{2}t^2 \overline{\boldsymbol{\zeta}}_{\mathbf{x}}) \tag{22}$$

By setting $t = 1$ and $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} = \overline{\log}_{\mathbf{x}}(\mathbf{y})$ (i.e., , $\langle \mathbf{x}, \mathbf{y}\rangle_q = r^2$ since $\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q = 0$), we have:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(1) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{1}{r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y} - \mathbf{x}\rangle_q \left(\mathbf{x} + \frac{1}{2}(\mathbf{y} - \mathbf{x})\right) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{1}{2r^2}\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y}\rangle_q (\mathbf{x} + \mathbf{y}) \tag{23}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q}{\langle \mathbf{x}, \mathbf{y}\rangle_q + r^2}(\mathbf{y} + \mathbf{x}) \tag{24}$$

(2) If $\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q > 0$, we find:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q}{r\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}} \left(\sin\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}\left(1 - \cos\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}{r}\right)\right)\overline{\boldsymbol{\zeta}}_{\mathbf{x}}\right)$$

By setting $t = 1$ and $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} = \overline{\log}_{\mathbf{x}}(\mathbf{y})$ (i.e., , $\langle \mathbf{x}, \mathbf{y}\rangle_q \in (-r^2, r^2)$), and using the fact that $\sin(\cos^{-1}(x)) = \sqrt{1 - x^2}$, we find:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(1) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y}\rangle_q}{r^2\sqrt{1 - (\frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})^2}} \left(\left(\sqrt{1 - (\frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})^2}\right)\mathbf{x} + \frac{1 - \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2}}{\sqrt{1 - (\frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})^2}}\left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2}\mathbf{x}\right)\right) \tag{25}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y}\rangle_q}{r^2(1 - (\frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})^2)} \left((1 - (\frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})^2)\mathbf{x} + (1 - \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})\left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2}\mathbf{x}\right)\right) \tag{26}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y}\rangle_q}{r^2(1 + \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2}))} \left((1 + \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2})\mathbf{x} + \left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y}\rangle_q}{r^2}\mathbf{x}\right)\right) \tag{27}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q}{\langle \mathbf{x}, \mathbf{y}\rangle_q + r^2}(\mathbf{y} + \mathbf{x}) \tag{28}$$

(3) If $\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q < 0$, we find:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(t) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q}{r\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}} \left(\sinh\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}\left(\cosh\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\zeta}}_{\mathbf{x}}, \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q|}}{r}\right) - 1\right)\overline{\boldsymbol{\zeta}}_{\mathbf{x}}\right)$$

By setting $t = 1$ and $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} = \overline{\log}_{\mathbf{x}}(\mathbf{y})$ (i.e., , $\langle \mathbf{x}, \mathbf{y} \rangle_q > r^2$), and using the fact that $\sinh(\cosh^{-1}(x)) = \sqrt{x^2 - 1}$, we find:

$$\overline{\boldsymbol{\xi}}_{\mathbf{x}}(1) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y} \rangle_q}{r^2 \sqrt{(\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2})^2 - 1}} \left(\left(\sqrt{(\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2})^2 - 1} \right) \mathbf{x} + \frac{\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2} - 1}{\sqrt{(\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2})^2 - 1}} \left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2} \mathbf{x} \right) \right) \tag{29}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y} \rangle_q}{r^2 ((\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2})^2 - 1)} \left(((\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2})^2 - 1) \mathbf{x} + (\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2} - 1) \left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2} \mathbf{x} \right) \right) \tag{30}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \mathbf{y} \rangle_q}{r^2 (1 + \frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2}))} \left((1 + \frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2}) \mathbf{x} + \left(\mathbf{y} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2} \mathbf{x} \right) \right) \tag{31}$$

$$= \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_q}{\langle \mathbf{x}, \mathbf{y} \rangle_q + r^2} (\mathbf{y} + \mathbf{x}) \tag{32}$$

In all cases, we define $P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) := \overline{\boldsymbol{\xi}}_{\mathbf{x}}(1)$ as formulated in Eq. (6). The parallel translation $P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}})$ is then performed along the **minimizing** geodesic $\overline{\gamma}_{\mathbf{x} \to \overline{\boldsymbol{\zeta}}_{\mathbf{x}}}$ defined such that $\overline{\gamma}_{\mathbf{x} \to \overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(1) = \mathbf{y}$ (i.e., $\overline{\exp}_{\mathbf{x}}^{-1}(\mathbf{y}) = \overline{\boldsymbol{\zeta}}_{\mathbf{x}}$) and $\overline{\gamma}'_{\mathbf{x} \to \overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(1) = P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}})$.

One can also verify that we have:

$$\forall \overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}} \mathcal{S}^{p,q}_r, \ P^{\overline{\gamma}}_{\mathbf{y} \frown \mathbf{x}} \left(P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) \right) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} \text{ if } \langle \mathbf{x}, \mathbf{y} \rangle_q > -r^2. \tag{33}$$

Nonexistence of (unbroken) geodesic joining pairs of points. $\mathbf{x} \in \mathcal{S}^{p,q}_r$ and $\mathbf{y} \in \mathcal{S}^{p,q}_r$ are joined by a geodesic iff $\langle \mathbf{x}, \mathbf{y} \rangle_q > -r^2$ or $\mathbf{y} = -\mathbf{x}$. A proof can be found in Appendix C.2 of [15] for the pseudo-hyperboloid $\mathcal{Q}^{q,p}_r$ that is anti-isometric to $\mathcal{S}^{p,q}_r$ as explained in Section B.4.

B.2 Indefinite elliptic space $\mathcal{P}^{p,q}_r$

The differential geometry tools of $\mathcal{P}^{p,q}_r$ depend on those of the pseudo-sphere described above. We recall that the canonical map $\pi : \mathcal{S}^{p,q}_r \to \mathcal{P}^{p,q}_r$ is defined as $\forall \mathbf{x} \in \mathcal{S}^{p,q}_r, \pi(\mathbf{x}) := [\mathbf{x}] = \{\mathbf{x}, -\mathbf{x}\}$.

B.2.1 Geodesic, exponential map and distance

Geodesic. By using the notation of the main paper, we recall that $\gamma = \pi \circ \overline{\gamma}$ and:

$$\forall \mathbf{x} \in \mathcal{S}^{p,q}_r, \boldsymbol{\xi} \in T_{[\mathbf{x}]} \mathcal{P}^{p,q}_r, \ \overline{\boldsymbol{\xi}}_{\mathbf{x}} = \text{lift}_{\mathbf{x}}(\boldsymbol{\xi}) = -\text{lift}_{-\mathbf{x}}(\boldsymbol{\xi}) = -\overline{\boldsymbol{\xi}}_{-\mathbf{x}} \tag{34}$$

and we have for all $t \in \mathbb{R}$, $\gamma_{[\mathbf{x}] \to \boldsymbol{\xi}}(t) = \{\overline{\gamma}_{\mathbf{x} \to \overline{\boldsymbol{\xi}}_{\mathbf{x}}}(t), \overline{\gamma}_{-\mathbf{x} \to \overline{\boldsymbol{\xi}}_{-\mathbf{x}}}(t)\}$.

Exponential map. The exponential map $\exp_{[\mathbf{x}]} : T_{[\mathbf{x}]} \mathcal{P}^{p,q}_r \to \mathcal{P}^{p,q}_r$ is defined such that

$$\exp_{[\mathbf{x}]}(\boldsymbol{\xi}) := \gamma_{[\mathbf{x}] \to \boldsymbol{\xi}}(1) = \{\overline{\gamma}_{\mathbf{x} \to \overline{\boldsymbol{\xi}}_{\mathbf{x}}}(1), \ \overline{\gamma}_{-\mathbf{x} \to \overline{\boldsymbol{\xi}}_{-\mathbf{x}}}(1)\} = [\overline{\exp}_{\mathbf{x}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}})]. \tag{35}$$

Logarithm map. $\log_{[\mathbf{x}]} := \exp_{[\mathbf{x}]}^{-1}$ is the inverse function of the exponential map. We can write:

$$\text{lift}_{\mathbf{x}} \left(\log_{[\mathbf{x}]}([\mathbf{y}]) \right) = \begin{cases} \overline{\log}_{\mathbf{x}}(\mathbf{y}) & \text{if } \langle \mathbf{x}, \mathbf{y} \rangle_q > 0 \\ \overline{\log}_{\mathbf{x}}(-\mathbf{y}) & \text{if } \langle \mathbf{x}, \mathbf{y} \rangle_q < 0 \end{cases} \tag{36}$$

where $\overline{\log}_{\mathbf{x}}$ is defined in Eq. (15). In theory, $\log_{[\mathbf{x}]}$ is not defined if $\langle \mathbf{x}, \mathbf{y} \rangle_q = 0$ because there exist two minimizing geodesics. In practice, we consider that its lift equals $\overline{\log}_{\mathbf{x}}(\mathbf{y})$ if $\langle \mathbf{x}, \mathbf{y} \rangle_q = 0$.

Geodesic distance. As stated in the paper, the geodesic distance $\mathsf{d}_\gamma(\cdot, \cdot)$ is then formulated:

$$\forall [\mathbf{x}] \in \mathcal{P}^{p,q}_r, [\mathbf{y}] \in \mathcal{P}^{p,q}_r, \ \mathsf{d}_\gamma([\mathbf{x}], [\mathbf{y}]) = \begin{cases} r \cosh^{-1}(|\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2}|) & \text{if } |\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2}| \geq 1 \\ r \cos^{-1} \ (|\frac{\langle \mathbf{x}, \mathbf{y} \rangle_q}{r^2}|) & \text{otherwise.} \end{cases} \tag{37}$$

It satisfies $\mathsf{d}_\gamma([\mathbf{x}], [\mathbf{y}]) = \min\{\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x}, \mathbf{y}), \overline{\mathsf{d}}_{\overline{\gamma}}(-\mathbf{x}, \mathbf{y})\}$.

B.2.2 $\mathcal{P}_r^{p,q}$ is a quotient manifold

We now explain why $\mathcal{P}_r^{p,q} := \mathcal{S}_r^{p,q}/\pm 1 = \mathcal{S}_r^{p,q}/\pm \mathbf{I}$ is a quotient manifold. The explanation is based on the "Orbit manifolds" section of Chapter 7 of [21] (see also page 192 of [21]). We first recall its Definition 6 and Proposition 7.

Definition B.2.1 (Definition 6 of Chapter 7 of [21])**.** A group Γ of diffeomorphisms of a manifold $\mathcal{M}$ is properly discontinuous (and acts freely) provided:

(PD1) Each point $\mathbf{x} \in \mathcal{M}$ has a neighborhood $\mathcal{A}$ such that if $\phi(\mathcal{A})$ meets $\mathcal{A}$ for $\phi \in \Gamma$ then $\phi =$id.

(PD2) Points $\mathbf{x}, \mathbf{z} \in \mathcal{M}$ not in the same orbit have neighborhoods $\mathcal{A}$ and $\mathcal{B}$ such that for every $\phi \in \Gamma$, $\phi(\mathcal{A})$ and $\mathcal{B}$ are disjoint.

Proposition B.2.1 (Proposition 7 of Chapter 7 of [21])**.** Let Γ be a properly discontinuous group of diffeomorphisms of a manifold $\mathcal{M}$. There is a unique way to make $\mathcal{M}/\Gamma$ a manifold so that the natural map $\pi : \mathcal{M} \to \mathcal{M}/\Gamma$ is a covering map.

In our case, we have $\mathcal{M} = \mathcal{S}_r^{p,q}$, and $\Gamma = \pm 1 = \pm \mathbf{I}$ is a group of diffeomorphisms of $\mathcal{S}_r^{p,q}$. To be more precise, Γ is composed of the identity map $\mathbf{x} \mapsto \mathbf{x}$ and the antipodal map $\mathbf{x} \mapsto -\mathbf{x}$. For all $\mathbf{x} \in \mathcal{S}_r^{p,q}$, the set $\{\phi(\mathbf{x}) : \phi \in \Gamma\} = [\mathbf{x}] = \{-\mathbf{x}, \mathbf{x}\}$ is called the orbit of $\mathbf{x}$ under Γ. The collection of all such orbits is our set $\mathcal{P}_r^{p,q} := \mathcal{S}_r^{p,q}/\pm 1$.

(PD1) is satisfied when the neighborhood $\mathcal{A}$ of $\mathbf{x} \in \mathcal{S}_r^{p,q}$ is defined as $\mathcal{A} = \{\mathbf{y} \in \mathcal{S}_r^{p,q} : \langle \mathbf{x}, \mathbf{y} \rangle_q > 0\}$.

(PD2) is satisfied when $\mathbf{z} \neq \pm\mathbf{x}$ by determining some neighborhood small enough for both $\mathbf{z}$ and $\pm\mathbf{x}$ so that they are disjoint.

By definition, each point $\mathbf{x} \in \mathcal{S}_r^{p,q}$ has a connected neighborhood $\mathcal{A} = \{\mathbf{y} \in \mathcal{S}_r^{p,q} : \langle \mathbf{x}, \mathbf{y} \rangle_q > 0\}$ that is evenly covered by π since it maps each component of $\pi^{-1}(\mathcal{A})$ diffeomorphically onto $\mathcal{V}$ (see Definition 7 of Chapter A of [21]). It is then a covering map and $\mathcal{P}_r^{p,q}$ is a quotient manifold.

It is worth noting that $\mathcal{P}_r^{p,q}$ is briefly mentioned in page 214 of [21]. It is also called an *indefinite elliptic space* and defined in Equation (12.2.2a) of [30], and the Riemannian case of elliptic geometry is briefly explained in page 74 of [30].

B.3 Pseudo-hyperboloid $\mathcal{Q}_r^{q,p}$

We recall here the differential geometry tools (from [15]) specific to the pseudo-hyperboloid which is defined as the following set: $\mathcal{Q}_r^{q,p} := \{\mathbf{x} \in \mathbb{R}^{q,p+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{p+1} = -r^2\}$.

Geodesic. The geodesic $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}} : \mathbb{R} \to \mathcal{Q}_r^{q,p}$ satisfying $\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(0) = \mathbf{x}$ and $\overline{\gamma}'_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(0) = \overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{Q}_r^{q,p}$ is formulated for all $t \in \mathbb{R}$:

$$\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(t) = \begin{cases} \cos\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\mathbf{x} \; + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}} \sin\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} < 0 \\ \mathbf{x} + t\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} = 0 \\ \cosh\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}} \sinh\left(\frac{t\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} > 0 \end{cases} \tag{38}$$

Exponential map. The exponential map $\overline{\exp}_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{Q}_r^{q,p} \to \mathcal{Q}_r^{q,p}$ is defined such that $\forall \overline{\boldsymbol{\xi}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{Q}_r^{q,p}, \overline{\exp}_{\mathbf{x}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) = \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(1)$. We then have:

$$\overline{\exp}_{\mathbf{x}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) = \begin{cases} \cos\left(\frac{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\mathbf{x} \; + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}} \sin\left(\frac{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} < 0 \\ \mathbf{x} + \overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} = 0 \\ \cosh\left(\frac{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\mathbf{x} + \frac{r}{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}} \sinh\left(\frac{\sqrt{|\langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1}|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} & \text{if } \langle \overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}} \rangle_{p+1} > 0 \end{cases} \tag{39}$$

Logarithm map. The logarithm map $\overline{\log}_{\mathbf{x}}$ is defined as the inverse of the exponential map $\overline{\exp}_{\mathbf{x}}$ on a normal neighborhood of $\mathbf{x} \in \mathcal{Q}_r^{q,p}$ denoted by $\mathcal{U}_{\mathbf{x}} = \{\mathbf{y} \in \mathcal{Q}_r^{q,p} : \frac{\langle \mathbf{x}, \mathbf{y} \rangle_{p+1}}{r^2} < 1\}$. It is then formulated:

$$\forall \mathbf{y} \in \mathcal{U}_{\mathbf{x}},\ \overline{\log}_{\mathbf{x}}(\mathbf{y}) = \begin{cases} \frac{\cosh^{-1}(-\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2})}{\sqrt{(\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2})^2-1}} \left(\mathbf{y} + \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2}\mathbf{x}\right) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2} < -1 \\ \mathbf{y} - \mathbf{x} & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2} = -1 \\ \frac{\cos^{-1}(-\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2})}{\sqrt{1-(\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2})^2}} \left(\mathbf{y} + \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2}\mathbf{x}\right) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2} \in (-1,1) \end{cases} \tag{40}$$

Geodesic "distance". The geodesic distance $\overline{\mathsf{d}}_{\overline{\gamma}} : \mathcal{Q}_r^{q,p} \times \mathcal{Q}_r^{q,p} \to \mathbb{R}$ is then:

$$\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) = \sqrt{|\langle \overline{\log}_{\mathbf{x}}(\mathbf{y}), \overline{\log}_{\mathbf{x}}(\mathbf{y})\rangle_{p+1}|} = \begin{cases} r\cosh^{-1}(-\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2}) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2} \leq -1 \\ r\cos^{-1}\ (-\frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2}) & \text{if } \frac{\langle \mathbf{x},\mathbf{y}\rangle_{p+1}}{r^2} \in (-1,1) \end{cases} \tag{41}$$

Parallel transport on $\mathcal{Q}_r^{q,p}$. The parallel transport connecting $\mathbf{x} \in \mathcal{Q}_r^{q,p}$ to $\mathbf{y} \in \mathcal{Q}_r^{q,p}$ is formulated:

$$P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{y}}(\overline{\boldsymbol{\xi}}_{\mathbf{x}}) := \overline{\boldsymbol{\xi}}_{\mathbf{x}} - \frac{\langle \mathbf{y}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_{p+1}}{\langle \mathbf{x},\mathbf{y}\rangle_{p+1} - r^2}(\mathbf{y}+\mathbf{x}) \qquad \text{where} \quad \langle \mathbf{x},\mathbf{y}\rangle_{p+1} < r^2 \tag{42}$$

B.4 Anti-isometry between the pseudo-sphere and the pseudo-hyperboloid

We now explain why the pseudo-sphere $\mathcal{S}_r^{p,q} := \left\{\mathbf{x} \in \mathbb{R}^{p+1,q} : \langle \mathbf{x},\mathbf{x}\rangle_q = r^2\right\}$ is anti-isometric to the pseudo-hyperboloid $\mathcal{Q}_r^{q,p} := \{\mathbf{x} \in \mathbb{R}^{q,p+1} : \langle \mathbf{x},\mathbf{x}\rangle_{p+1} = -r^2\}$. This can actually be generalized to the anti-isometry between $\mathbb{R}^{p+1,q}$ and $\mathbb{R}^{q,p+1}$ that we describe below.

Let us note the vectors $\mathbf{x} = (x_0, x_1, \dots, x_{d-1}, x_d)^\top \in \mathbb{R}^{p+1,q}$ and $\mathbf{y} = (y_0, y_1, \dots, y_{d-1}, y_d)^\top \in \mathbb{R}^{p+1,q}$. We can construct vectors in $\mathbb{R}^{q,p+1}$ that reverse the order of the elements of $\mathbf{x}$ and $\mathbf{y}$. We obtain the following vectors $\mathbf{a} = (x_d, x_{d-1}, \dots, x_1, x_0)^\top \in \mathbb{R}^{q,p+1}$ and $\mathbf{b} = (y_d, y_{d-1}, \dots, y_1, y_0)^\top \in \mathbb{R}^{q,p+1}$. By definition of our scalar product in Eq. (1), the anti-isometry between $\mathbb{R}^{p+1,q}$ and $\mathbb{R}^{q,p+1}$ corresponds to:

$$\langle \mathbf{x},\mathbf{y}\rangle_q = -\langle \mathbf{a},\mathbf{b}\rangle_{p+1}. \tag{43}$$

For instance, for the hyperboloid, let us assume that $\mathbf{x} = (x_0, x_1, \dots, x_{d-1}, x_d)^\top \in \mathcal{S}_1^{0,q}$ and we note $\mathbf{a} = (x_d, x_{d-1}, \dots, x_1, x_0)^\top \in \mathcal{Q}_1^{q,0}$. We find:

$$\langle \mathbf{a},\mathbf{a}\rangle_1 = -\langle \mathbf{x},\mathbf{x}\rangle_q = -x_0^2 + \sum_{j=1}^{d} x_j^2 = -1. \tag{44}$$

B.5 Explanation of Figure 1

We give the definition of space-like and time-like geodesics in Appendix B.1. We recall that $r = 1$ in the figure.

Space-like geodesic. In Figure 1, $\mathbf{x}$ and $\mathbf{y}$ are connected by a space-like geodesic. Therefore, according to Eq. (16), the geodesic distance between $\mathbf{x}$ and $\mathbf{y}$ is $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{y}) = r\cos^{-1}\ (\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2})$ and the geodesic distance between $[\mathbf{x}]$ and $[\mathbf{y}]$ is $\mathsf{d}_\gamma([\mathbf{x}],[\mathbf{y}]) = r\cos^{-1}\ (|\frac{\langle \mathbf{x},\mathbf{y}\rangle_q}{r^2}|) = \overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},-\mathbf{y})$.

Time-like geodesic. In Figure 1, $\mathbf{x}$ and $\mathbf{z}$ are connected by a time-like geodesic. Therefore, the geodesic distance between $\mathbf{x}$ and $\mathbf{z}$ is $\overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{z}) = r\cosh^{-1}\ (\frac{\langle \mathbf{x},\mathbf{z}\rangle_q}{r^2})$ and the geodesic distance between $[\mathbf{x}]$ and $[\mathbf{z}]$ is $\mathsf{d}_\gamma([\mathbf{x}],[\mathbf{z}]) = r\cos^{-1}\ (|\frac{\langle \mathbf{x},\mathbf{z}\rangle_q}{r^2}|) = \overline{\mathsf{d}}_{\overline{\gamma}}(\mathbf{x},\mathbf{z})$.

Null geodesic. For completeness, the geodesic "distance" between two points joined by a null geodesic is 0 even if the two points are distinct.

B.6 Hyperbolic and elliptic parts of the ultrahyperbolic manifold

In the main paper, we state that $\mathcal{P}_r^{p,q}$ contains hyperbolic and elliptic parts. Our explanation is similar to the one in [15].

• **Elliptic parts.** We first recall that if all the time dimensions of $\mathcal{P}_r^{p,q}$ are set to 0, then the considered manifold can be written $\mathcal{P}_r^{p,0} \times \{\mathbf{0}\}$ which corresponds to elliptic geometry.

Moreover, in spherical geometry, geodesics are all written in the following way:

$$\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(t) = \cos\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle|}}{r}\right)\mathbf{x} \; + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle|}} \sin\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} \tag{45}$$

Their formulation is then very similar to the formulation of our space-like geodesics of Eq. (13) except that a different scalar product is used. In fact, it corresponds to a special case of our scalar product when the number of time dimensions is zero.

• **Hyperbolic parts.** We also recall that if all the space dimensions except one of $\mathcal{P}_r^{p,q}$ are set to 0, then the considered manifold is diffeomorphic to $\{\mathbf{0}\} \times \mathcal{P}_r^{0,q}$ which corresponds to the hyperboloid model of hyperbolic geometry.

Moreover, in the hyperboloid model of hyperbolic geometry, geodesics are all written:

$$\overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\xi}}_{\mathbf{x}}}(t) = \cosh\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\mathbf{x} \; + \frac{r}{\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}} \sinh\left(\frac{t\sqrt{|\langle\overline{\boldsymbol{\xi}}_{\mathbf{x}}, \overline{\boldsymbol{\xi}}_{\mathbf{x}}\rangle_q|}}{r}\right)\overline{\boldsymbol{\xi}}_{\mathbf{x}} \tag{46}$$

Their formulation is then similar to the formulation of our time-like geodesics of Eq. (13) except that a larger number of time dimensions is used in our case.

In conclusion, our proposed geometry is more general and manages to describe relationships considered in elliptic and hyperbolic geometries.

C Descent direction and optimization

C.1 Descent direction of Section 3.4

Proof. We provide here the detailed proof that the negative of $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} \in \mathcal{H}_{\mathbf{p}}$ is a descent direction. We recall that $\mathbf{x} := \overline{\exp}_{\mathbf{p}}\left(\varphi_\theta(\boldsymbol{x})\right) \in \mathcal{S}_r^{p,q}$ and $\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} := \mathrm{lift}_{\mathbf{p}}\left(P^{\gamma}_{[\mathbf{x}]\frown[\mathbf{p}]}(Df([\mathbf{x}]))\right) \in T_{\mathbf{p}}\mathcal{S}_r^{p,q}$.

• We first consider the case where $\langle \mathbf{x}, \mathbf{p}\rangle_q \geq 0$.

Let us consider some tangent vector $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ and some point $\mathbf{y} \in \mathcal{S}_r^{p,q}$ defined such that $\overline{f}(\mathbf{y}) = \overline{f} \circ \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(1)$. By exploiting Taylor's first-order approximation, the function $\overline{f} \circ \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}$ can be approximated at $t = 1$ by:

$$\overline{f}(\mathbf{y}) = \overline{f} \circ \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(1) \simeq \overline{f} \circ \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}}(0) + (\overline{f} \circ \overline{\gamma}_{\mathbf{x}\to\overline{\boldsymbol{\zeta}}_{\mathbf{x}}})'(0) = \overline{f}(\mathbf{x}) + \langle D\overline{f}(\mathbf{x}), \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q \tag{47}$$

where $D\overline{f}(\mathbf{x}) \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ is the pseudo-Riemannian gradient of $\overline{f}$ at $\mathbf{x}$ (see Section 4.2 of [15] for details).

Our goal is to determine some tangent vector $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ such that it is a descent direction. In other words, we want $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} \in T_{\mathbf{x}}\mathcal{S}_r^{p,q}$ to satisfy $\overline{f}(\mathbf{y}) < \overline{f}(\mathbf{x})$ (i.e., $\langle D\overline{f}(\mathbf{x}), \overline{\boldsymbol{\zeta}}_{\mathbf{x}}\rangle_q < 0$).

We also recall that our neural network φ_θ maps to $T_{\mathbf{p}}\mathcal{S}_r^{p,q}$ but $\overline{\boldsymbol{\zeta}}_{\mathbf{x}}$ lies in $T_{\mathbf{x}}\mathcal{S}_r^{p,q}$, which is a different tangent space if $\mathbf{p} \neq \mathbf{x}$. The parallel transport allows us to work with both tangent spaces. To simplify the notation, we define the following tangent vector:

$$\bar{\boldsymbol{\chi}}_{\mathbf{p}} := \boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} = \mathrm{lift}_{\mathbf{p}}\left(P^{\gamma}_{[\mathbf{x}]\frown[\mathbf{p}]}(Df([\mathbf{x}]))\right) = P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{p}}(D\overline{f}(\mathbf{x})) \in T_{\mathbf{p}}\mathcal{S}_r^{p,q}. \tag{48}$$

As explained in Section B.1.2, the geodesic $\overline{\gamma}_{\mathbf{p}\to\bar{\boldsymbol{\chi}}_{\mathbf{p}}}$ satisfies the properties $\overline{\gamma}_{\mathbf{p}\to\bar{\boldsymbol{\chi}}_{\mathbf{p}}}(0) = \mathbf{p}$, $\overline{\gamma}'_{\mathbf{p}\to\bar{\boldsymbol{\chi}}_{\mathbf{p}}}(0) = \bar{\boldsymbol{\chi}}_{\mathbf{p}}$ and $\overline{\gamma}'_{\mathbf{p}\to\bar{\boldsymbol{\chi}}_{\mathbf{p}}}(1) = D\overline{f}(\mathbf{x})$. We then have $P^{\overline{\gamma}}_{\mathbf{p}\frown\mathbf{x}}(\bar{\boldsymbol{\chi}}_{\mathbf{p}}) = D\overline{f}(\mathbf{x})$.

It is worth noting that we also have the following property: $\forall \overline{\boldsymbol{\xi}}_{\mathbf{p}} \in T_{\mathbf{p}}\mathcal{S}_r^{p,q}, \mathbf{G}\overline{\boldsymbol{\xi}}_{\mathbf{p}} \in T_{\mathbf{p}}\mathcal{S}_r^{p,q}$. Let us then define $\overline{\boldsymbol{\zeta}}_{\mathbf{x}}$ such that $\overline{\boldsymbol{\zeta}}_{\mathbf{x}} := P^{\overline{\gamma}}_{\mathbf{p}\frown\mathbf{x}}(-\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}) = P^{\overline{\gamma}}_{\mathbf{p}\frown\mathbf{x}}(-\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}})$. From Eq. (33), we know that $-\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}} = P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{p}}(\overline{\boldsymbol{\zeta}}_{\mathbf{x}})$ and $\bar{\boldsymbol{\chi}}_{\mathbf{p}} = P^{\overline{\gamma}}_{\mathbf{x}\frown\mathbf{p}}(D\overline{f}(\mathbf{x}))$.

Due to the linear isometry property of the parallel transport (see page 66 of [21]), we have:

$$\langle D\overline{f}(\mathbf{x}), \overline{\boldsymbol{\zeta}}_{\mathbf{x}} \rangle_q = \langle P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{p}}(D\overline{f}(\mathbf{x})), P^{\overline{\gamma}}_{\mathbf{x} \frown \mathbf{p}}(-\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}}) \rangle_q = \langle \bar{\boldsymbol{\chi}}_{\mathbf{p}}, -\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}} \rangle_q = -\|\bar{\boldsymbol{\chi}}_{\mathbf{p}}\|^2 \leq 0 \tag{49}$$

where $\|\cdot\|$ denotes the standard Euclidean norm defined as $\forall \mathbf{x}, \|\mathbf{x}\| := \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. Eq. (49) is zero iff $\bar{\boldsymbol{\chi}}_{\mathbf{p}} = \mathbf{0}$, and negative otherwise. It is also worth noting that $\bar{\boldsymbol{\chi}}_{\mathbf{p}} = \mathbf{0}$ iff $D\overline{f}(\mathbf{x}) = \mathbf{0}$ (i.e., $\mathbf{x}$ is a stationary point). This shows that the negative of $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$ is a descent direction.

Due to the properties of the exponential map, the differential of the exponential map $\mathrm{d}(\overline{\exp}_{\mathbf{p}})_{\mathbf{0}}$ at the origin $\mathbf{0}$ satisfies the following property:

$$\langle \mathrm{d}(\overline{\exp}_{\mathbf{p}})_{\mathbf{0}}(\bar{\boldsymbol{\chi}}_{\mathbf{p}}), -\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}} \rangle_q = \langle \bar{\boldsymbol{\chi}}_{\mathbf{p}}, -\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}} \rangle_q = -\|\bar{\boldsymbol{\chi}}_{\mathbf{p}}\|^2 \leq 0 \tag{50}$$

Eq. (50) implies that $-\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$ is a descent direction of the neural network $\varphi_\theta : \mathcal{X} \to T_{\mathbf{p}}\mathcal{S}_r^{p,q}$.

• The case where $\langle \mathbf{x}, \mathbf{p} \rangle_q < 0$ is similar to the case above except that we now have:

$$\overline{f}(\mathbf{y}) = \overline{f} \circ \overline{\gamma}_{-\mathbf{x} \to \overline{\boldsymbol{\zeta}}_{-\mathbf{x}}}(1) \simeq \overline{f}(-\mathbf{x}) + (\overline{f} \circ \overline{\gamma}_{-\mathbf{x} \to \overline{\boldsymbol{\zeta}}_{-\mathbf{x}}})'(0) = \overline{f}(-\mathbf{x}) + \langle D\overline{f}(-\mathbf{x}), \overline{\boldsymbol{\zeta}}_{-\mathbf{x}} \rangle_q \tag{51}$$

$$= \overline{f}(\mathbf{x}) + \langle -D\overline{f}(\mathbf{x}), \overline{\boldsymbol{\zeta}}_{-\mathbf{x}} \rangle_q \tag{52}$$

where $\overline{\boldsymbol{\zeta}}_{-\mathbf{x}} := P^{\overline{\gamma}}_{\mathbf{p} \frown -\mathbf{x}}(-\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}) \in T_{-\mathbf{x}}\mathcal{S}_r^{p,q}$ and $D\overline{f}(-\mathbf{x}) = -D\overline{f}(\mathbf{x}) \in T_{-\mathbf{x}}\mathcal{S}_r^{p,q}$.

We also have $\bar{\boldsymbol{\chi}}_{\mathbf{p}} := \boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} = \mathrm{lift}_{\mathbf{p}}\left(P^{\gamma}_{[\mathbf{x}] \frown [\mathbf{p}]}(Df([\mathbf{x}]))\right) = P^{\overline{\gamma}}_{-\mathbf{x} \frown \mathbf{p}}(-D\overline{f}(\mathbf{x})) \in T_{\mathbf{p}}\mathcal{S}_r^{p,q}$, which implies $\langle -D\overline{f}(\mathbf{x}), \overline{\boldsymbol{\zeta}}_{-\mathbf{x}} \rangle_q = \langle \bar{\boldsymbol{\chi}}_{\mathbf{p}}, -\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}} \rangle_q = \langle \bar{\boldsymbol{\chi}}_{\mathbf{p}}, -\mathbf{G}\bar{\boldsymbol{\chi}}_{\mathbf{p}} \rangle_q = -\|\bar{\boldsymbol{\chi}}_{\mathbf{p}}\|^2 \leq 0$.

This completes the proof. $\square$

C.2 Optimizing the MLP in the toy experiment

In the toy experiment, we define a new PyTorch autograd function to define the exponential map $\overline{\exp}_{\mathbf{p}}$ as explained in Section 3.4. The custom gradient of our autograd function is $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$.

Naïvely using (the negative of) $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$ as descent direction and exploiting standard backpropagation decreases the optimized function because all the hidden layers already lie in some space equipped with a positive definite metric tensor [9].

C.3 Optimizing the Graph Convolutional Network

For the GCN introduced in Section 4, the optimized parameters are the matrices $\mathbf{W}^k$. To be fair with the baselines, we modified the code of Liu *et al.* [18] that is available at the following address: `https://github.com/facebookresearch/hgnn`

We added a Python class for our ultrahyperbolic manifold, it is very similar to the Lorentz manifold Python class. Our code replaces the standard Lorentz inner product (that corresponds to our scalar product in the special case where $p = 0$) used in [18] with our scalar product, its induced exponential/logarithm map and geodesic distance. We also modified the activation function as explained in Section 4. Standard backpropagation is used to train the parameters $\mathbf{W}^k$ that exploit operations over the horizontal space of the positive pole $\mathbf{p}$ as explained in the paper. To have a fair comparison, we used the optimizer of [18] and did not use the optimizer introduced in Section 3.4.

D Experiments

D.1 Type of resources used and amount of compute

We ran all our experiments on Zachary's karate club dataset and the node classification task on a machine equipped with a 6-core Intel i7-7800X CPU and NVIDIA GeForce RTX 3090 GPU. The machine was also used to run most of our graph classification experiments.

Since the Reddit-multi-12K dataset requires more than 24GB of VRAM, we ran each experiment of the Reddit dataset on a single 32 GB NVIDIA Tesla V100 GPU of an NVIDIA DGX-1 server. Most experiments take several minutes. Some graph classification experiments take several hours and the longest experiment (one split of Reddit-multi-12k) takes one day.

D.2 Zachary's karate club dataset

Parameters and hyperparameters. We train our framework as explained in the main paper and Section C.2. In practice, we define a new PyTorch autograd function to define the exponential map $\overline{\exp}_{\mathbf{p}}$. The custom gradient of our autograd function is $\mathbf{G}\boldsymbol{\lambda}_{[\mathbf{x}],\mathbf{p}}$. Concerning the choice of hyperparameters (e.g., temperature τ, optimizer and learning rate) , we chose the same hyperparameter values as Law & Stam [15]. During training, we use a standard Stochastic Gradient Descent (SGD) optimizer without momentum, with learning rate of 10^{-7}. We run our experiments for 25,000 iterations (which are also epochs since the dataset is small). We chose a standard MLP with 3 hidden layers to show that our optimizer can be used with neural networks. Other architectures can be used.

Figure 2 (right) and other illustrative two-dimensional plots. As a qualitative way to understand the method, we plot two-dimensional representations that were learned for different runs. For all the illustrative figures, we replace the geodesic distance used in Eq. (12) by the squared geodesic distance. It tends to give nicer illustrations.

- We plot two-dimensional projections of learned points lying on the non-Riemannian manifold $\mathcal{P}_1^{1,1}$ in Fig. 2 (right) and Fig 3 (see caption of the figure for details). These projections lie in non-Euclidean space so they should not be interpreted by using standard Euclidean distances. Instead, the figures on the right correspond to spacetime diagrams. As explained in the caption of Fig. 3, points lying on an oblique line have very small distance. Nonetheless, we can see a clear separation between nodes of different factions.

- We plot the same kind of two-dimensional hyperbolic and elliptic representations in Fig. 5 and Fig 4, respectively. Although the separation between factions is clear, the learned node representations do not satisfy the standard structure of a tree or a cycle graph. For instance, high-level nodes of the hierarchy (i.e., nodes v_1 and v_{34}) do not lie closer to the origin than low-level nodes although this is generally the case when hyperbolic representations are used to learn trees [19, 20].

Evaluation metrics. Following the evaluation protocol of [15], we take the capacity matrix $\mathbf{C} \in \mathbb{R}^{n \times n}$ of [33] which defines the level of friendship between the different members. We then consider instead its symmetrized version $\mathbf{S} = \mathbf{C} + \mathbf{C}^\top$. The score $s_i = \sum_{j=1}^{n} \mathbf{S}_{ij}$ defines the importance of the node v_i in the hierarchy. The higher the score, the more important the node is in the hierarchy.

These s_i scores are then used to calculate the Spearman's rank correlation coefficient between the selected s_i scores (top 5 or top 10) and corresponding δ_i scores. As reported in Table 1, ultrahyperbolic representations are more correlated with the node importance in the hierarchy.

Other proxy to quantify importance in hyperbolic space. In machine learning, when hyperbolic embeddings are used to represent hierarchies or trees, a standard way to determine the importance of nodes is to compare the Euclidean norm of the embeddings in the Poincaré ball (or equivalently on the hyperboloid) [19, 20]. High-level nodes tend to have smaller Euclidean norm in hyperbolic geometry. In the first column of Table 6, we report the different scores when the Euclidean norm of the learned hyperbolic representations is used as a proxy of the importance. The second column corresponds to the scores reported in Table 1 of the main paper (i.e., sum of the δ_i scores).

According to the results in Table 6, the ℓ_2-norm is a worse indicator of importance than δ_i scores for this dataset due to the presence of cycles in the graph. This observation is also in accordance with the qualitative two-dimensional results of Fig. 4 where nodes v_1 and v_{34} do not lie closer to the origin than other nodes.

Table 6: Evaluation scores for the different learned representations (mean $\pm$ standard deviation)

Evaluation metric	Hyperbolic with ℓ_2 norm as proxy	Hyperbolic with δ_i score as proxy
Rank of first leader	$\mathbf{2.2 \pm 1.0}$	2.5 ± 0.7
Rank of second leader	7.3 ± 2.4	$\mathbf{3.8 \pm 1.0}$
top 5 Spearman's ρ	0.30 ± 0.44	$\mathbf{0.36 \pm 0.22}$
top 10 Spearman's ρ	0.22 ± 0.21	$\mathbf{0.38 \pm 0.18}$

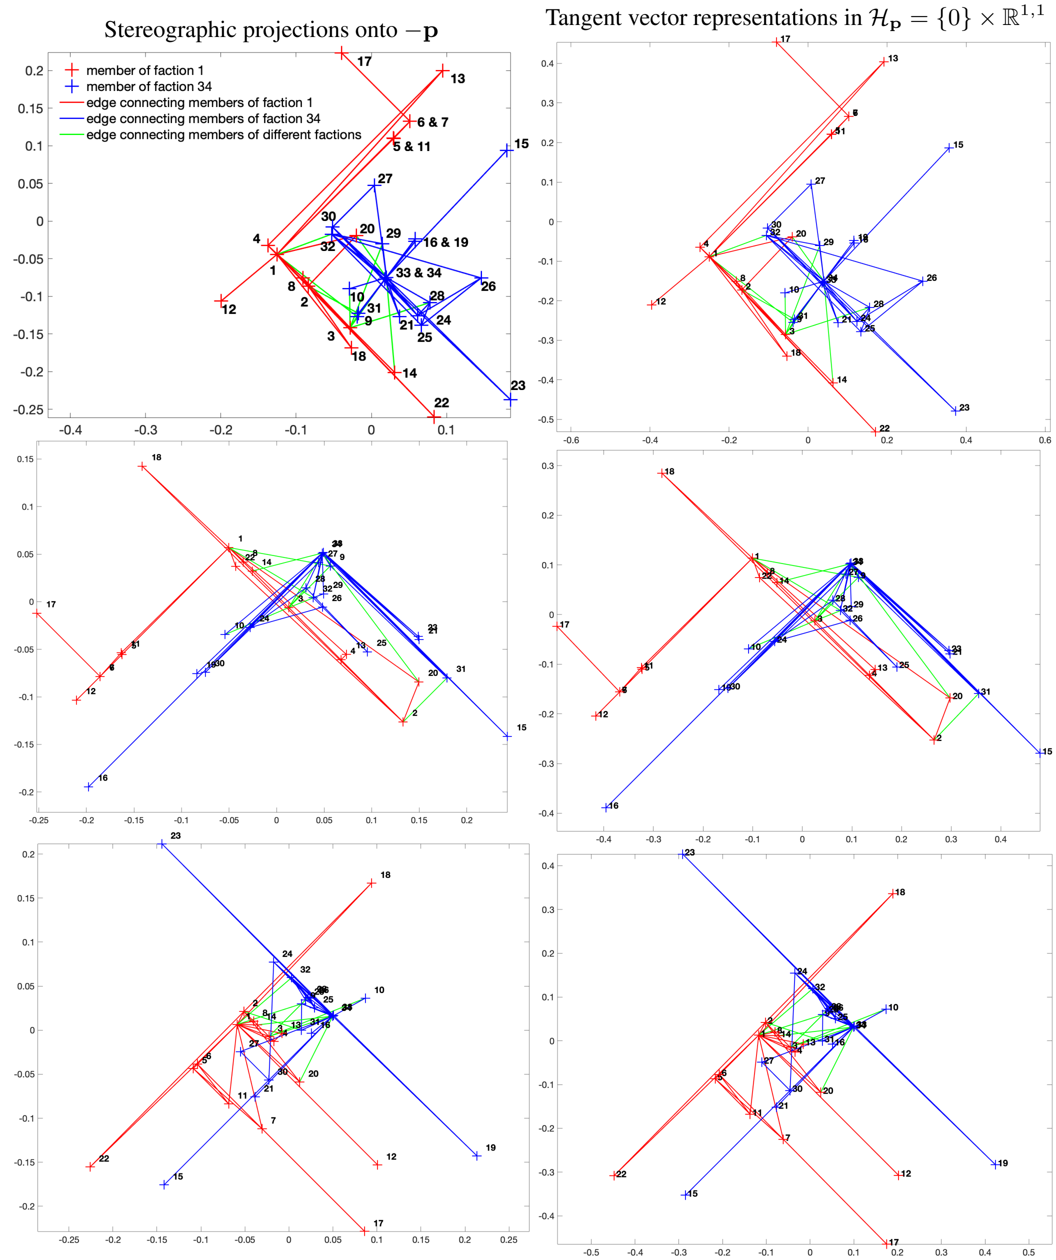

Figure 3: (left) Stereographic projection of learned node representations in $\mathcal{P}_1^{1,1}$ for three different initializations. (right) Tangent vector representations of node representations. For every node representation $[\mathbf{x}_i] \in \mathcal{P}_1^{1,1}$, we plot the last two elements of its tangent vector representation: $\boldsymbol{\xi}_i = \text{lift}_{\mathbf{p}}\left(\log_{[\mathbf{p}]}([\mathbf{x}_i])\right) \in \mathcal{H}_{\mathbf{p}} = \{0\} \times \mathbb{R}^{1,1}$. Tangent vector representations are easier to interpret since they lie in some space diffeomorphic to $\mathbb{R}^{1,1}$. Let us consider two vectors $\mathbf{a} = (a_1, a_2) \in \mathbb{R}^{1,1}$ and $\mathbf{b} = (b_1, b_2) \in \mathbb{R}^{1,1}$. Their distance in $\mathbb{R}^{1,1}$ is $\sqrt{|\langle \mathbf{a} - \mathbf{b}, \mathbf{a} - \mathbf{b} \rangle_q|} = \sqrt{|(a_1 - b_1)^2 - (a_2 - b_2)^2|}$, which explains why similar examples (i.e., connected by an edge) are joined by an oblique line. Their distance in that space is very small and does not follow the intuition of the standard Euclidean distance.

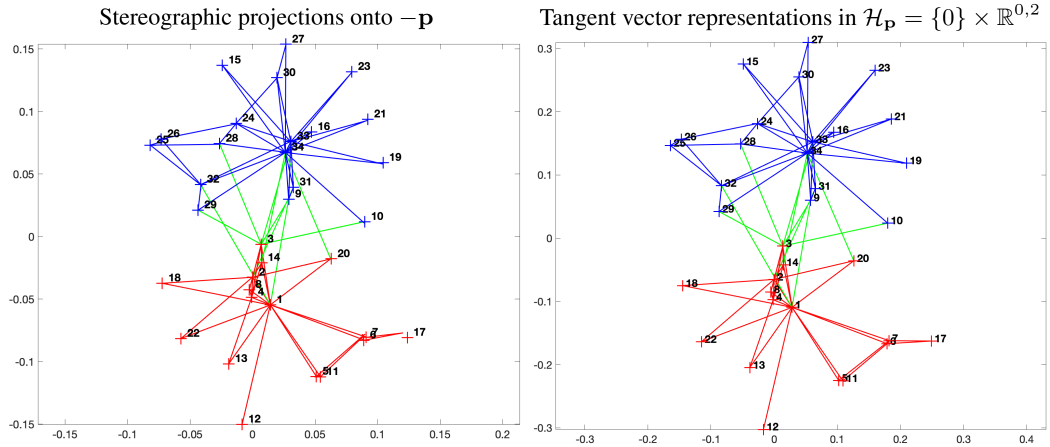

Figure 4: (left) Stereographic projection of learned **hyperbolic** node representations in $\mathcal{P}_1^{0,2}$. In the machine learning literature, they are also called **Poincaré representations**. (right) Tangent vector representations of node representations. For every node representation $[\mathbf{x}_i] \in \mathcal{P}_1^{1,1}$, we plot the last two elements of its tangent vector representation: $\boldsymbol{\xi}_i = \text{lift}_{\mathbf{p}}\left(\log_{[\mathbf{p}]}([\mathbf{x}_i])\right) \in \mathcal{H}_{\mathbf{p}} = \{0\} \times \mathbb{R}^{0,2}$. It is worth noting that, since the represented graph is not a tree, the high-level nodes (i.e., nodes v_1 and v_{34}) do not have smaller Euclidean norm than other nodes in the hierarchy.

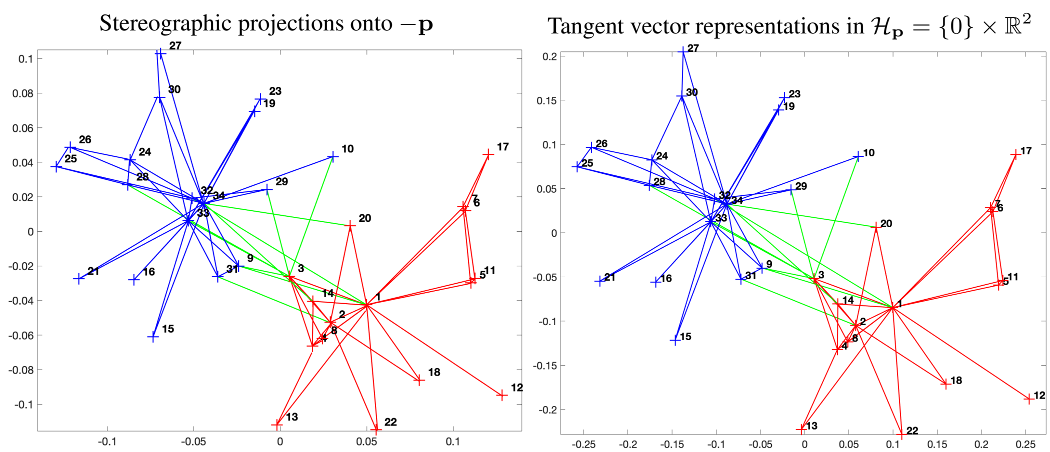

Figure 5: (left) Stereographic projection of learned **elliptic** node representations in $\mathcal{P}_1^{2,0}$. (right) Tangent vector representations of node representations. For every node representation $[\mathbf{x}_i] \in \mathcal{P}_1^{2,0}$, we plot the last two elements of its tangent vector representation: $\boldsymbol{\xi}_i = \text{lift}_{\mathbf{p}}\left(\log_{[\mathbf{p}]}([\mathbf{x}_i])\right) \in \mathcal{H}_{\mathbf{p}} = \{0\} \times \mathbb{R}^2$.

Table 7: Test node classification accuracy with 10-dimensional manifolds

Dataset	$\mathbb{R}^{10}$ (Euclidean)	$\mathcal{P}_1^{0,10}$ (Hyperbolic)	$\mathcal{P}_1^{1,9}$	$\mathcal{P}_1^{2,8}$	$\mathcal{P}_1^{9,1}$	$\mathcal{P}_1^{10,0}$ (Elliptic)
Citeseer	58.4 ± 2.1	56.4 ± 2.9	$\mathbf{62.2 \pm 2.1}$	60.9 ± 2.8	60.4 ± 3.4	61.3 ± 2.7
Cora	67.8 ± 4.8	72.6 ± 2.1	$\mathbf{75.1 \pm 1.6}$	73.7 ± 2.3	73.3 ± 2.7	71.9 ± 1.9
Pubmed	73.1 ± 2.5	$\mathbf{75.3 \pm 1.6}$	74.9 ± 1.9	75.0 ± 1.0	75.1 ± 1.3	$\mathbf{75.3 \pm 0.8}$

Table 8: Test node classification accuracy with 600-dimensional manifolds

Dataset	$\mathbb{R}^{600}$	$\mathcal{P}_1^{0,600}$	$\mathcal{P}_1^{1,599}$
Citeseer	70.9 ± 0.4	70.8 ± 0.4	70.6 ± 0.5
Cora	81.6 ± 0.4	81.9 ± 0.3	82.0 ± 0.4
Pubmed	79.0 ± 0.5	79.0 ± 0.8	78.9 ± 0.8

D.3 Node and graph classification

We now give details about the experiments of Section 5.2. As explained in Section C.3, for the node and graph classification tasks, we simply adapted the code of Liu *et al.* [18] to the ultrahyperbolic case. We refer the reader to [18] for more details since our experimental protocol is the same.

Data preprocessing and choice of splits. To download the datasets, we used the splits extracted from Liu's project page (`https://github.com/facebookresearch/hgnn`). The node classification extraction script is `download_node.sh` and the graph classification extraction script is `data_preprocess.py` which provides 10 fixed splits per dataset to perform 10-fold cross validation.

Prototype-based classification. Following Section 3 of [18], the output of an ultrahyperbolic neural network with K steps is a set of node representations in ultrahyperbolic space: $\{\mathbf{h}_1^K, \dots, \mathbf{h}_{|V|}^K\}$ where each $\mathbf{h}_i^K$ lies on the manifold. A list of prototypes (called "centroids" in [18]) is created $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_{|\mathcal{C}|}\}$ where each $\mathbf{c}_j$ lies on the same manifold as $\mathbf{h}_i^K$. All the prototypes are points, they are learned jointly with the GNN using backpropagation.

A distance matrix $\mathbf{D} \in \mathbb{R}^{|V| \times |\mathcal{C}|}$ defined such that $\mathbf{D}_{ij} = \mathsf{d}(\mathbf{h}_i^K, \mathbf{c}_j)$ is created. In practice, d is the geodesic distance. It satisfies the properties in Section 3.3 and our optimization framework can be used.

• **Node classification.** Let us note C the number of node classes and $\mathbf{W} \in \mathbb{R}^{|\mathcal{C}| \times C}$ some matrix to be learned. The posterior probability distribution to determine the category of each node is calculated as follows: $\mathbf{Y} = \text{softmax}(\mathbf{DW})$ where the j-th element of the i-th row of $\mathbf{Y}$ corresponds to the probability that the i-th node belongs to the j-th category. Cross-entropy is used for learning.

• **Graph classification.** For graph-level predictions, average pooling is first used to combine the distances of different nodes into a single score per node. As done in [18], a fully connected layer is then used with standard cross-entropy to perform graph classification.

Choice of parameters and hyperparameters. In the same way as Section D.2 for Zachary's karate club dataset, our code is based on the code of Liu *et al.* [18] as explained in Section C.3. To be fair with the baselines, we take the parameters available at `https://github.com/facebookresearch/hgnn/tree/master/params` that were used for the hyperbolic manifold. We only replace the hyperboloid by $\mathcal{P}_r^{p,q}$, and we adapt the activation function as explained in Section 4.

For instance, for node classification, we use the following parameters: `https://github.com/facebookresearch/hgnn/blob/master/params/NodeClassificationHyperbolicParams.py` (i.e., same optimizer, learning rate, number of prototypes, number of layers etc). We do the same thing for the graph classification task.

Reported results. In the tables of results, the baselines "Euclidean", "Poincaré " and "Lorentz" correspond to the implementations of [18]. Liu *et al.* show that their implementation matches the scores of the standard GCN. We did not manage to reproduce their results for the collab and reddit datasets even when we tried different optimizers, learning rates, activation functions, number of centroids. We then reran their code and reported the obtained results in the main paper. We report results when the manifold is 10-dimensional (resp. 600-dimensional) in Table 7 (resp. Table 8).

Table 9: Evaluation scores for the different learned representations on Zachary's karate club dataset (mean $\pm$ standard deviation)

Manifold	Distance	Rank of first leader	Rank of second leader	top 5 Spearman's ρ	top 10 Spearman's ρ
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	d_{ℓ_1}	1.8 ± 0.5	3.4 ± 0.7	0.47 ± 0.25	0.52 ± 0.13
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	d_{ℓ_2}	1.9 ± 0.8	3.4 ± 0.9	0.47 ± 0.20	0.51 ± 0.18
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	$\mathsf{d}_{\min}$	3.0 ± 2.3	7.2 ± 3.4	0.23 ± 0.23	0.39 ± 0.15
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	d_{ℓ_1}	2.2 ± 0.7	3.8 ± 0.7	0.24 ± 0.29	0.48 ± 0.17
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	d_{ℓ_2}	2.0 ± 0.7	3.6 ± 1.5	0.48 ± 0.24	0.50 ± 0.23
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	$\mathsf{d}_{\min}$	3.6 ± 2.5	8.0 ± 3.6	0.16 ± 0.30	0.48 ± 0.24
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	d_{ℓ_1}	1.8 ± 0.7	3.4 ± 0.8	0.48 ± 0.19	0.51 ± 0.17
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	d_{ℓ_2}	1.8 ± 0.7	3.6 ± 0.9	0.31 ± 0.21	0.52 ± 0.16
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	$\mathsf{d}_{\min}$	3.0 ± 2.3	7.8 ± 3.2	0.13 ± 0.42	0.46 ± 0.22

D.4 Comparison with products of Riemannian space forms

In the main paper, we do not compare $\mathcal{P}_r^{p,q}$ to products of spherical and hyperbolic manifolds [3, 11] because these product manifolds do not have constant curvature and we could similarly consider products of pseudo-sphere of same dimension to add more complexity, which would have made the paper hard to read. In this subsection, we report these comparisons.

Notation. $\mathbb{S}^p_{r_1} := \mathcal{S}^{p,0}_{r_1}$ denotes the p-sphere of radius r_1 (embedded in a $(p+1)$-dimensional Euclidean space). Similarly, $\mathbb{H}^q_{r_2}$ denotes the q-dimensional hyperboloid of "radius" r_2 and embedded in a $(q+1)$-dimensional space. Following [3, 11], the radii $r_1 > 0$ and $r_2 > 0$ are trained parameters (both initialized at 1) and we define the following distance metrics for the product manifold $\mathbb{S}^p_{r_1} \times \mathbb{H}^q_{r_2}$ (see [11] for details):

- The geodesic ℓ_2 distance: $\mathsf{d}_{\ell_2}((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) := \sqrt{\mathsf{d}_1^2(\mathbf{x}_1, \mathbf{x}_2) + \mathsf{d}_2^2(\mathbf{y}_1, \mathbf{y}_2)}$
- The ℓ_1 distance: $\mathsf{d}_{\ell_1}((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) := \mathsf{d}_1(\mathbf{x}_1, \mathbf{x}_2) + \mathsf{d}_2(\mathbf{y}_1, \mathbf{y}_2)$
- The min distance: $\mathsf{d}_{\min}((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) := \min(\mathsf{d}_1(\mathbf{x}_1, \mathbf{x}_2), \mathsf{d}_2(\mathbf{y}_1, \mathbf{y}_2))$

where $\mathsf{d}_1(\mathbf{x}_1, \mathbf{x}_2) := r_1 \cos^{-1}(\langle \mathbf{x}_1, \mathbf{x}_2 \rangle / r_1^2)$ and $\mathsf{d}_2(\mathbf{y}_1, \mathbf{y}_2) := r_2 \cosh^{-1}(|\langle \mathbf{y}_1, \mathbf{y}_2 \rangle_q / r_2^2|)$ are the geodesic distances of the p-sphere of radius r_1 and $q-$hyperboloid of radius r_2, respectively.

D.4.1 Zachary's karate club dataset

We report the scores on Zachary's karate club dataset in Table 9 by using the following evaluation metrics: Rank of first leader, Rank of second leader, Spearman's ρ for the top 5 nodes, Spearman's ρ for the top 10 nodes. These evaluation metrics quantify how much the chosen distance extracts the hierarchy information in the graph.

All these product manifolds perform better than Riemannian space forms (see Table 1) but worse than the quotient manifold $\mathcal{P}_1^{p,q}$. It is worth noting that the best performing distance metrics are d_{ℓ_1} and d_{ℓ_2}. They both add the spherical and hyperbolic distances and then explicitly enforce both a spherical and hyperbolic structure when comparing pairs of samples. The fact that they perform worse than the geodesic distance of $\mathcal{P}_r^{p,q}$ indicates that explicitly constructing hyperbolic and spherical parts to the manifold by using products of Riemannian manifolds may not be optimal depending on the selected pairs.

Interestingly, the distance metric $\mathsf{d}_{\min}$ that selects some hyperbolic or spherical distance depending on the pair of samples performs much worse. This is in contrast with our approach that also intrinsically selects a elliptic or hyperbolic type of distance depending on the pair of compared samples (see Eq. (7)). However, the selection in Eq. (7) is based on the (intrinsic) geodesic "distance" of the manifold $\mathcal{P}_r^{p,q}$. Experimental results suggest that the fact that $\mathcal{P}_r^{p,q}$ intrinsically contains hyperbolic and elliptic parts due to the indefiniteness of the metric tensor allows us to better describe hierarchical relationships between samples when the hierarchical graph contains cycles.

Table 10: Evaluation scores for the learned 4-dimensional representations in the node classification task (mean $\pm$ standard deviation)

Manifold	Distance	Citeseer	Cora	Pubmed
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	d_{ℓ_1}	43.4 ± 2.6	56.6 ± 2.9	68.5 ± 4.8
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	d_{ℓ_2}	46.8 ± 2.1	57.6 ± 2.4	71.5 ± 2.1
$\mathbb{S}^1_{r_1} \times \mathbb{H}^3_{r_2}$	$\mathsf{d}_{\min}$	40.7 ± 3.9	47.5 ± 2.5	63.0 ± 1.4
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	d_{ℓ_1}	45.9 ± 1.9	60.4 ± 2.8	70.5 ± 2.6
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	d_{ℓ_2}	47.2 ± 2.1	60.5 ± 3.2	71.1 ± 2.5
$\mathbb{S}^2_{r_1} \times \mathbb{H}^2_{r_2}$	$\mathsf{d}_{\min}$	44.4 ± 2.3	55.2 ± 4.9	70.1 ± 2.1
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	d_{ℓ_1}	47.3 ± 2.0	56.5 ± 2.4	71.9 ± 2.1
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	d_{ℓ_2}	48.1 ± 2.1	60.8 ± 2.8	72.5 ± 1.8
$\mathbb{S}^3_{r_1} \times \mathbb{H}^1_{r_2}$	$\mathsf{d}_{\min}$	43.6 ± 3.2	55.2 ± 2.9	68.9 ± 2.6

D.4.2 Results in node classification

We ran the same kind of experiment as above in the node classification task described in Section 5.2. We report in Table 10 the results obtained with 4-dimensional manifolds and the same distance metrics (see Table 3 for comparison).

Once again, $\mathsf{d}_{\min}$ performs worse than the other distance metrics that perform slight better than hyperbolic and elliptic distances but are still outperformed by our proposed distances on the Cora and Citeseer datasets.

We ran similar experiments for 10-dimensional manifolds. The conclusion is similar.