

Low-rank approximations of hyperbolic embeddings

Pratik Jawanpuria, Mayank Meghwanshi, and Bamdev Mishra

Abstract—The hyperbolic manifold is a smooth manifold of negative constant curvature. While the hyperbolic manifold is well-studied in the literature, it has gained interest in the machine learning and natural language processing communities lately due to its usefulness in modeling continuous hierarchies. Tasks with hierarchical structures are ubiquitous in those fields and there is a general interest to learning hyperbolic representations or embeddings of such tasks. Additionally, these embeddings of related tasks may also share a low-rank subspace. In this work, we propose to learn hyperbolic embeddings such that they also lie in a low-dimensional subspace. In particular, we consider the problem of learning a low-rank factorization of hyperbolic embeddings. We cast these problems as manifold optimization problems and propose computationally efficient algorithms. Empirical results illustrate the efficacy of the proposed approach.

I. INTRODUCTION

Learning hyperbolic representation of entities have gained recent interest in the machine learning community [1], [2], [3], [4], [5]. In particular, hyperbolic embeddings have been shown to be well-suited for various natural language processing problems [3], [6], [7] that require modeling hierarchical structures such as knowledge graphs, hypernymy hierarchies, organization hierarchy, and taxonomies, among others. The reason being learning representations in the hyperbolic space provides a principled approach for integrating structural information encoded in such (discrete) entities into continuous space.

The hyperbolic space is a non-Euclidean space and has constant negative curvature. The latter property enables it to grow exponentially even in dimension as low as two. Hence, the hyperbolic space has been considered to model trees and complex networks, among others [8], [9]. Figure 1(a) is an example of representing a part of mammal taxonomy tree in a hyperbolic space (two-dimensional Poincaré ball). Hyperbolic embeddings (numerical representations of tasks) have been considered in several applications such as question answering system [10], recommender systems [11], [12], link prediction [13], [14], natural language inference [15], vertex classification [16], and machine translation [17].

In this paper, we consider the setting in which an additional low-rank structure may also exist among the learned hyperbolic embeddings. Such a setting may arise when the hierarchical entities are closely related. We propose to learn a low-rank approximation of the given (high dimensional) hyperbolic embeddings. Conceptually, we model high dimensional hyperbolic embeddings as a product of a low-dimensional subspace and low-dimensional hyperbolic em-

beddings. The optimization problem is cast on the product of the Stiefel and hyperbolic manifolds. We develop an efficient Riemannian trust-region algorithm for solving it. We evaluate the proposed approach on real-world datasets: on the problem of reconstructing taxonomy hierarchies from the embeddings. We observe that the performance of proposed approach match the original embeddings even in low-rank settings.

The outline of the paper is as follows. Section II discusses two popular models of representing hyperbolic space in the Euclidean setting. In Section III, we present our formulation to approximate given hyperbolic embeddings in a low-rank setting. The optimization algorithm is discussed in Section IV. The experimental results are presented in Section V and Section VI concludes the paper.

II. BACKGROUND

In this section, we briefly discuss the basic concepts of hyperbolic geometry. Interested readers may refer [18], [19] for more details.

The hyperbolic space (of dimension ≥ 2) is a Riemannian manifold with a constant negative sectional curvature. Similar to the Euclidean or spherical spaces, it is isotropic. However, the Euclidean space is flat (zero curvature) and the spherical space is positively curved. As a result of negative curvature, the circumference and area of a circle in hyperbolic space grow exponentially with the radius. In contrast, the circumference and area of a circle in the hyperbolic space grow linearly and quadratically, respectively, in Euclidean setting. Hence, hyperbolic spaces *expand* faster than the Euclidean spaces. Informally, hyperbolic spaces may be viewed as a continuous counterpart to discrete trees as the metric properties of a two-dimensional hyperbolic space and a b -ary tree (a tree with branching factor b) are similar. Hence, trees can be embedded into a two-dimensional hyperbolic space while keeping the overall distortion arbitrarily small. In contrast, Euclidean spaces cannot attain this result even with unbounded number of dimensions.

Since hyperbolic models cannot be represented within Euclidean space without distortion, several (equivalent) models exist for representing hyperbolic spaces for computation purpose. The models are conformal to the Euclidean space and points in one model can be transformed to be represented in another model, while preserving geometric properties such as isometry. However, no model captures all the properties of the hyperbolic geometry. Two hyperbolic models, in particular, have received much interest recently in the machine learning community: the Poincaré ball model and the hyperboloid model.

The authors are with Microsoft, Hyderabad, India. Emails: {prjawanp, namegh, bamdevm}@microsoft.com.

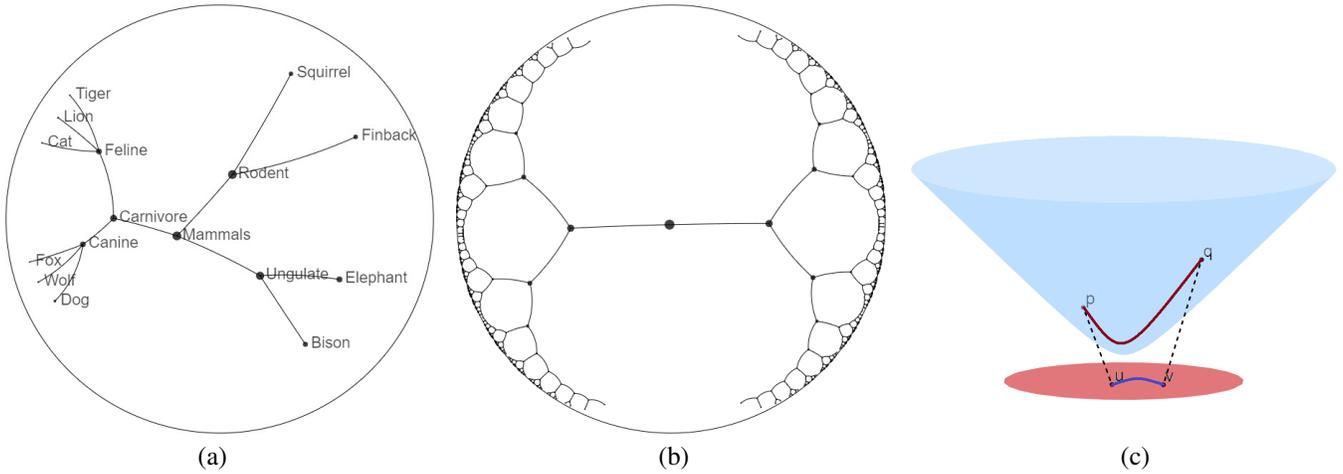


Fig. 1. (a) An example of the hyperbolic space (2-dimensional Poincaré ball model \mathbb{B}^2) being used to represent a mammal taxonomy. This taxonomy is a part of WordNet [20]; (b) A tree is embedded in \mathbb{B}^2 . The two subtrees from the root are regular trees. All the edges have the same hyperbolic length, computed using (1); (c) The Poincaré disk (\mathbb{B}^2) may be viewed as a stereoscopic projection of the hyperboloid model (\mathbb{H}^2). Points p and q lie on \mathbb{H}^2 and points u and v are their projections, respectively, onto the \mathbb{B}^2 . The maroon curve \mathbb{H}^2 is the geodesic between p and q , which projects to the blue geodesic path between u and v on \mathbb{B}^2 . Figure best viewed in color.

A. Poincaré ball model

The Poincaré ball is a n -dimensional hyperbolic space defined as the interior of the n -dimensional unit (Euclidean) ball:

$$\mathbb{B}^n = \{\mathbf{u} \in \mathbb{R}^n \mid \|\mathbf{u}\| < 1\},$$

where $\|\cdot\|$ denotes the Euclidean norm. The distance between two points $\mathbf{u}, \mathbf{v} \in \mathbb{B}^n$ in the Poincaré ball model is given by

$$d_{\mathbb{B}}(\mathbf{u}, \mathbf{v}) = \operatorname{arccosh} \left(1 + 2 \frac{\|\mathbf{u} - \mathbf{v}\|^2}{(1 - \|\mathbf{u}\|^2)(1 - \|\mathbf{v}\|^2)} \right) \quad (1)$$

and the Poincaré norm is given by

$$\|\mathbf{u}\|_{\mathbb{B}} := d_{\mathbb{B}}(\mathbf{0}, \mathbf{u}) = 2 \operatorname{arctanh}(\|\mathbf{u}\|).$$

We observe that the distance between a pair of points near the boundary of the Poincaré ball (Euclidean norm close to unity) grows much faster than distance between the points close to the center (Euclidean norm close to zero). In addition, the distance within the Poincaré ball varies smoothly with respect to points \mathbf{u} and \mathbf{v} . These properties are helpful embedding discrete hierarchical structures such as trees in hyperbolic spaces and obtain continuous embeddings which respect the tree metric structure. For instance, the origin of the Poincaré ball may be mapped to the root node of the tree as the root node is relatively closer to all other nodes (points). The leaf nodes can be placed near the boundary to ensure they are relatively distant from other leaf nodes. Additionally, the shortest path between a pair of points is usually via a point closer to the origin, just as the shortest path between two nodes in a tree is via their parent nodes. Figure 1(b) shows a Poincaré disk (\mathbb{B}^2) embedding a tree with two regular subtrees.

B. Hyperboloid model

Let $\bar{\mathbf{u}}, \bar{\mathbf{v}} \in \mathbb{R}^{n+1}$ such that $\bar{\mathbf{u}} = \begin{bmatrix} u_0 \\ \mathbf{u} \end{bmatrix}$ and $\bar{\mathbf{v}} = \begin{bmatrix} v_0 \\ \mathbf{v} \end{bmatrix}$ and $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$. The Lorentz scalar product $\langle \cdot, \cdot \rangle_{\mathcal{L}}$ of two vectors

$\bar{\mathbf{u}}$ and $\bar{\mathbf{v}}$ is defined as

$$\langle \bar{\mathbf{u}}, \bar{\mathbf{v}} \rangle_{\mathcal{L}} = \bar{\mathbf{u}}^{\top} \mathbf{L} \bar{\mathbf{v}} = -u_0 v_0 + \mathbf{u}^{\top} \mathbf{v}, \quad (2)$$

where \mathbf{L} is a $(n+1)$ -dimensional diagonal matrix

$$\mathbf{L} = \begin{bmatrix} -1 & \mathbf{0}_n^{\top} \\ \mathbf{0}_n & \mathbf{I}_n \end{bmatrix},$$

$\mathbf{0}_n$ is the n -dimensional zero column vector, and \mathbf{I}_n is the n -dimensional identity matrix.

The hyperboloid model, also known as the Lorentz model of hyperbolic geometry, is given by

$$\mathbb{H}^n = \{\bar{\mathbf{u}} \in \mathbb{R}^{n+1} \mid \langle \bar{\mathbf{u}}, \bar{\mathbf{u}} \rangle_{\mathcal{L}} = -1, u_0 > 0\}.$$

The model represents the upper sheet of an n -dimensional hyperboloid. From the constraint set, it can be observed that if $\bar{\mathbf{u}} \in \mathbb{H}^n$, then $u_0 = \sqrt{1 + \mathbf{u}^{\top} \mathbf{u}}$.

The distance between two points $\bar{\mathbf{u}}, \bar{\mathbf{v}} \in \mathbb{H}^n$ in the hyperboloid model is given by

$$d_{\mathbb{H}}(\bar{\mathbf{u}}, \bar{\mathbf{v}}) = \operatorname{arccosh}(-\langle \bar{\mathbf{u}}, \bar{\mathbf{v}} \rangle_{\mathcal{L}}). \quad (3)$$

As stated earlier, both the Poincaré ball and the hyperboloid models are equivalent and a mapping exists from one model to another [21]. Points on the hyperboloid can be mapped to the Poincaré ball by

$$h : \mathbb{H}^n \rightarrow \mathbb{B}^n, h(\bar{\mathbf{u}}) = \frac{\mathbf{u}}{u_0 + 1}, \bar{\mathbf{u}} \in \mathbb{H}^n.$$

The reverse mapping, $h^{-1} : \mathbb{B}^n \rightarrow \mathbb{H}^n$ is defined as follows:

$$h^{-1}(\mathbf{w}) = \frac{1}{(1 - \|\mathbf{w}\|^2)} \begin{bmatrix} (1 + \|\mathbf{w}\|^2) \\ 2\mathbf{w} \end{bmatrix}, \mathbf{w} \in \mathbb{B}^n.$$

Figure 1(c) shows a two-dimensional hyperboloid model \mathbb{H}^2 . It can be observed that the Poincaré disk \mathbb{B}^2 is obtained as a stereoscopic projection of \mathbb{H}^2 .

III. LOW-RANK PARAMETERIZATION IN HYPERBOLIC SPACE

As discussed earlier, hyperbolic embeddings are typically suitable for representing elements of hierarchical structures such as nodes of trees [3] and complex networks [8] to name a few. When the task involves closely related hierarchical concepts, additional low-rank structure may also exist among such hyperbolic embeddings. In this section, we propose a novel low-rank parameterization for hyperbolic embeddings. It should be noted that, unlike the Euclidean embeddings, incorporating a low-rank structure in the hyperbolic framework is non-trivial because of the hyperboloid constraints.

Let $\bar{\mathbf{X}}$ be a $(n+1) \times m$ matrix whose columns represent n -dimensional hyperbolic embeddings corresponding to m elements from a given hierarchical structure. For notational convenience, we represent $\bar{\mathbf{X}}$ and its i -th column $\bar{\mathbf{x}}_i$ as follows:

$$\bar{\mathbf{X}} = \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{X} \end{bmatrix} \text{ and } \bar{\mathbf{x}}_i = \begin{bmatrix} x_{0i} \\ \mathbf{x}_i \end{bmatrix}.$$

We propose to approximate $\bar{\mathbf{x}}_i$ as a low-dimensional hyperbolic embedding such that \mathbf{x}_i shares a latent low-dimensional subspace with \mathbf{x}_j (corresponding to $\bar{\mathbf{x}}_j$), for all $i, j = 1, \dots, m$. Mathematically, we propose the following $(r+1)$ -rank approximation for $\bar{\mathbf{x}}_i$:

$$\bar{\mathbf{x}}_i = \begin{bmatrix} x_{0i} \\ \mathbf{x}_i \end{bmatrix} \approx \begin{bmatrix} z_{0i} \\ \mathbf{U}\mathbf{z}_i \end{bmatrix} = \hat{\mathbf{z}}_i \quad \forall i = 1, \dots, m,$$

where $\hat{\mathbf{z}}_i \in \mathbb{H}^n$, $\mathbf{z}_i \in \mathbb{R}^r$, $\mathbf{U} \in \mathbb{R}^{n \times r}$, and $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$. We discuss below the consequences of the proposed model.

Firstly, we obtain $\bar{\mathbf{z}}_i = \begin{bmatrix} z_{0i} \\ \mathbf{z}_i \end{bmatrix} \in \mathbb{H}^r$. This is because $\hat{\mathbf{z}}_i \in \mathbb{H}^n$ implies

$$-z_{0i}^2 + (\mathbf{U}\mathbf{z}_i)^\top (\mathbf{U}\mathbf{z}_i) = -1.$$

$\bar{\mathbf{z}}_i \in \mathbb{H}^r$ follows from the above equality as $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$. Secondly, the matrix \mathbf{X} (corresponding to $\bar{\mathbf{X}}$) is modeled as a low-rank matrix as we approximate $\bar{\mathbf{X}}$ as $\mathbf{U}\mathbf{Z}$, where $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_m]$. Thirdly, the space complexity of embeddings reduces from $O(nm)$ (for $\bar{\mathbf{X}}$) to $O(nr + mr)$ (for \mathbf{U}, \mathbf{Z} and \mathbf{z}_0).

We propose to learn the proposed low-rank parameterization of $\bar{\mathbf{X}}$ by solving the optimization problem:

$$\begin{aligned} & \min_{\substack{\mathbf{U} \in \mathbb{R}^{n \times r}, \\ \bar{\mathbf{z}}_i \in \mathbb{H}^r \forall i}} \underbrace{\sum_{i=1}^m \ell(\bar{\mathbf{x}}_i, \hat{\mathbf{z}}_i)}_{f(\mathbf{U}, \bar{\mathbf{Z}}; \bar{\mathbf{X}})} \\ & \text{subject to } \mathbf{U}^\top \mathbf{U} = \mathbf{I}_r, \bar{\mathbf{z}}_i \in \mathbb{H}^r, \\ & \hat{\mathbf{z}}_i = \begin{bmatrix} 1 & \mathbf{0}_r^\top \\ \mathbf{0}_n & \mathbf{U} \end{bmatrix} \bar{\mathbf{z}}_i \quad \forall i = 1, \dots, m, \end{aligned} \quad (4)$$

where $\ell : \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} \rightarrow [0, \infty)$ is a *loss function* that measures the quality of the proposed approximation. Let function f denote the objective function in (4), i.e., $f(\mathbf{U}, \bar{\mathbf{Z}}) = \sum_{i=1}^m \ell(\bar{\mathbf{x}}_i, \hat{\mathbf{z}}_i)$.

We discuss the following three choices of f :

- 1) $f(\mathbf{U}, \bar{\mathbf{Z}}; \bar{\mathbf{X}}) = \|\mathbf{X} - \mathbf{U}\mathbf{Z}\|_F^2$:
we penalize the Euclidean distance between $\bar{\mathbf{X}}$ and $\mathbf{U}\mathbf{Z}$. This is because \mathbf{x}_0 and \mathbf{z}_0 are determined from the hyperboloid constraint given $\bar{\mathbf{X}}$ and $\mathbf{U}\mathbf{Z}$, respectively. We obtain a closed-form solution of (4) with this loss function and the solution involves computing a rank- r singular value decomposition of $\bar{\mathbf{X}}$. In Section V, we denote this approach by the term Method-1.
- 2) $f(\mathbf{U}, \bar{\mathbf{Z}}; \bar{\mathbf{X}}) = \|\bar{\mathbf{X}} - \hat{\mathbf{Z}}\|_F^2$:
we penalize the Euclidean distance between the (full) hyperbolic embeddings (matrices) $\bar{\mathbf{X}}$ and $\hat{\mathbf{Z}} = \begin{bmatrix} 1 & \mathbf{0}_r^\top \\ \mathbf{0}_n & \mathbf{U} \end{bmatrix} \bar{\mathbf{Z}}$. This approach is denoted by the term Method-2 in Section V.
- 3) $f(\mathbf{U}, \bar{\mathbf{Z}}; \bar{\mathbf{X}}) = \sum_i \text{arccosh}(-\langle \bar{\mathbf{x}}_i, \hat{\mathbf{z}}_i \rangle_{\mathcal{L}})^2$:
since the columns of $\bar{\mathbf{X}}$ and $\hat{\mathbf{Z}}$ are hyperbolic embeddings, we penalize the hyperbolic distance (3) between the corresponding embeddings. We denote it by the term Method-3.

In the next section, we propose a Riemannian trust-region algorithm for solving (4) with the loss function discussed in options 2) and 3) above.

IV. OPTIMIZATION

It should be noted that the variable \mathbf{U} in (4) belongs to the Stiefel manifold $\text{St}(r, n) := \{\mathbf{U} \in \mathbb{R}^{n \times r} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}_r\}$ [22] and the variable $\bar{\mathbf{z}}_i$ belongs to the r -dimensional hyperbolic manifold $\mathbb{H}^r := \{\bar{\mathbf{z}}_i \in \mathbb{R}^{r+1} : -z_{0i}^2 + \mathbf{z}_i^\top \mathbf{z}_i = -1, z_{0i} > 0\}$ for all $i = \{1, \dots, m\}$. Consequently, the constraint set of the proposed optimization problem (4) is a smooth manifold $\mathcal{M} := \text{St}(r, n) \times \mathbb{H}^r \times \dots \times \mathbb{H}^r$, which is the Cartesian product of the Stiefel and m hyperbolic manifolds of dimension r . The problem (4), therefore, now boils down to the manifold optimization problem:

$$\min_{y \in \mathcal{M}} f(y), \quad (5)$$

where y has the representation $y := (\mathbf{U}, \bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_m)$ and $f : \mathcal{M} \rightarrow \mathbb{R} : y \mapsto f(y) = \sum_{i=1}^m \ell(\bar{\mathbf{x}}_i, \hat{\mathbf{z}}_i)$ is a smooth function.

We tackle the problem (5) in the Riemannian optimization framework that translates it into an unconstrained optimization problem over the nonlinear manifold \mathcal{M} , now endowed with a Riemannian geometry [23]. In particular, the Riemannian geometry on manifolds imposes a metric (inner product) structure on \mathcal{M} , which in turn allows to generalize notions like the shortest distance between points (on the manifold) or the translation of vectors on manifolds. Following this framework many of the standard nonlinear optimization algorithms in the Euclidean space, e.g., steepest descent and trust-regions, generalize well to Riemannian manifolds in a systematic manner. The Riemannian framework allows to develop computationally efficient algorithms on manifolds [23].

Both the Stiefel and hyperbolic manifolds are Riemannian manifolds, and their geometries have been individually well-studied in the literature [13], [23]. Subsequently, the manifold of interest \mathcal{M} also has a Riemannian structure.

Below we list some of the basic optimization-related notions that are required to solve (5) with the Riemannian *trust-region* algorithm that exploits second-order information. The development of those notions follow the general treatment of manifold optimization discussed in [23, Chapter 7]. The Stiefel manifold related expressions follow from [23]. The hyperbolic related expressions follow from [13].

A. Metric and tangent space notions

Optimization on \mathcal{M} is worked out on the tangent space, which is the linearization of \mathcal{M} at a specific point. It is a vector space associated with each element of the manifold.

As \mathcal{M} is a product space, its tangent space is also the product space of the tangent spaces of the Stiefel $\text{St}(r, n)$ and hyperbolic \mathbb{H}^r manifolds. The characterization of the tangent space has the form:

$$\begin{aligned} T_y \mathcal{M} &:= T_{\mathbf{U}} \text{St}(r, n) \times T_{\bar{\mathbf{z}}_1} \mathbb{H}^r \times \dots \times T_{\bar{\mathbf{z}}_m} \mathbb{H}^r \\ &= \{(\xi_{\mathbf{U}}, \xi_{\bar{\mathbf{z}}_1}, \dots, \xi_{\bar{\mathbf{z}}_m}) : \text{symm}(\mathbf{U}^\top \xi_{\mathbf{U}}) = \mathbf{0}_r \text{ and} \\ &\quad \langle \bar{\mathbf{z}}_i, \xi_{\bar{\mathbf{z}}_i} \rangle_{\mathcal{L}} = 0 \text{ for all } i\}, \end{aligned} \quad (6)$$

where symm extracts the symmetric part of a matrix.

As discussed above, to impose a Riemannian structure on \mathcal{M} , a smooth metric (inner product) definition is required at each element of the manifold. A natural choice of the metric $g_y : T_y \mathcal{M} \times T_y \mathcal{M} \rightarrow \mathbb{R} : (\xi_y, \eta_y) \mapsto g_y(\xi_y, \eta_y)$ on \mathcal{M} is the summation of the individual Riemannian metrics on the Stiefel and hyperbolic manifolds. More precisely, we have

$$g_y(\xi_y, \eta_y) := \langle \xi_{\mathbf{U}}, \eta_{\mathbf{U}} \rangle + \sum_{i=1}^m \langle \xi_{\bar{\mathbf{z}}_i}, \eta_{\bar{\mathbf{z}}_i} \rangle_{\mathcal{L}}, \quad (7)$$

where $\xi_y = (\xi_{\mathbf{U}}, \xi_{\bar{\mathbf{z}}_1}, \dots, \xi_{\bar{\mathbf{z}}_m})$, $\eta_y = (\eta_{\mathbf{U}}, \eta_{\bar{\mathbf{z}}_1}, \dots, \eta_{\bar{\mathbf{z}}_m})$, $\langle \cdot, \cdot \rangle$ is the standard inner product, and $\langle \cdot, \cdot \rangle_{\mathcal{L}}$ is the Lorentz inner product (2).

It should be emphasized that the metric in (7) endows the manifold \mathcal{M} with a Riemannian structure and allows to develop various other notions of optimization in a straightforward manner.

One important ingredient required in optimization is the notion of an orthogonal projection operator Π_y from the space $\mathbb{R}^{n \times r} \times \mathbb{R}^{n+1} \dots \times \mathbb{R}^{n+1}$ to the tangent space $T_y \mathcal{M}$. Exploiting the product and Riemannian structure of \mathcal{M} , the projection operator characterization is obtained as the Cartesian product of the individual tangent space projection operator on the Stiefel and hyperbolic manifolds, both of which are well known. Specifically, if $(\zeta_{\mathbf{U}}, \zeta_{\bar{\mathbf{z}}_1}, \dots, \zeta_{\bar{\mathbf{z}}_m}) \in \mathbb{R}^{n \times r} \times \mathbb{R}^{n+1} \dots \times \mathbb{R}^{n+1}$, then its projection onto the tangent space $T_y \mathcal{M}$ is given by

$$\begin{aligned} \Pi_y(\zeta_{\mathbf{U}}, \zeta_{\bar{\mathbf{z}}_1}, \dots, \zeta_{\bar{\mathbf{z}}_m}) &:= (\zeta_{\mathbf{U}} - \mathbf{U} \text{symm}(\mathbf{U}^\top \zeta_{\mathbf{U}}), \\ &\quad \zeta_{\bar{\mathbf{z}}_1} + \bar{\mathbf{z}}_1 \langle \mathbf{z}_1, \zeta_{\bar{\mathbf{z}}_1} \rangle_{\mathcal{L}}, \\ &\quad \dots, \\ &\quad \zeta_{\bar{\mathbf{z}}_m} + \bar{\mathbf{z}}_m \langle \mathbf{z}_m, \zeta_{\bar{\mathbf{z}}_m} \rangle_{\mathcal{L}}). \end{aligned} \quad (8)$$

B. Retraction

An optimization algorithm on manifold requires computation of search direction and then following along it. While the computation of the search direction follows from the notions in Section IV-A, in this section we develop the

notion of “moving” along a search direction on the manifold. This is characterized by the *retraction* operation, which is the generalization of the the exponential map (that follows the geodesic) on the manifold. The retraction operator R_y takes in a tangent vector at $T_y \mathcal{M}$ and outputs an element on the manifold by approximating the geodesic [23, Definition 4.1.1].

Exploiting the product space of \mathcal{M} , a natural expression of the retraction operator is obtained by the Cartesian product of the individual retraction operations on the Stiefel and hyperbolic manifolds. If $\xi_y \in T_y \mathcal{M}$, then the retraction operation is given by

$$\begin{aligned} R_y(\xi_y) &:= (\text{uf}(\mathbf{U} + \xi_{\mathbf{U}}), \\ &\quad \bar{\mathbf{z}}_1 \cosh(\|\xi_{\bar{\mathbf{z}}_1}\|_{\mathcal{L}}) + \xi_{\bar{\mathbf{z}}_1} \sinh(\|\xi_{\bar{\mathbf{z}}_1}\|_{\mathcal{L}}) / \|\xi_{\bar{\mathbf{z}}_1}\|_{\mathcal{L}}, \\ &\quad \dots, \\ &\quad \bar{\mathbf{z}}_m \cosh(\|\xi_{\bar{\mathbf{z}}_m}\|_{\mathcal{L}}) + \xi_{\bar{\mathbf{z}}_m} \sinh(\|\xi_{\bar{\mathbf{z}}_m}\|_{\mathcal{L}}) / \|\xi_{\bar{\mathbf{z}}_m}\|_{\mathcal{L}}), \end{aligned} \quad (9)$$

where $\xi_y = (\xi_{\mathbf{U}}, \xi_{\bar{\mathbf{z}}_1}, \dots, \xi_{\bar{\mathbf{z}}_m})$, $\|\xi_{\bar{\mathbf{z}}_i}\|_{\mathcal{L}} = \sqrt{\langle \xi_{\bar{\mathbf{z}}_i}, \xi_{\bar{\mathbf{z}}_i} \rangle_{\mathcal{L}}}$ for all i , $\text{uf}(\cdot)$ extracts the orthogonal factor of a matrix, *i.e.*, $\text{uf}(\mathbf{A}) = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1/2}$.

C. Riemannian gradient and Hessian computations

Finally, we require the expressions of the Riemannian gradient and Hessian of f on \mathcal{M} . To this end, we first compute the derivatives of f in the Euclidean space. Let $\nabla_y f$ be the first derivative of f and its Euclidean directional derivative along ξ_y is $\text{D}\nabla_y f[\xi_y]$. The expressions of the partial derivatives for the squared Euclidean distance based loss functions mentioned in Section III are straightforward to compute. When the loss function is based on the squared hyperbolic distance (3), the expressions for $\nabla_y f$ and $\text{D}\nabla_y[\xi_y]$ are discussed in [24].

Once the partial derivatives of f are known, converting them to their Riemannian counterparts on \mathcal{M} follows from the theory of Riemannian optimization [23, Chapter 3]. The expressions are

$$\text{Riemannian gradient } \text{grad}_y f = \Pi_y(\nabla_y f), \text{ and} \quad (10)$$

$$\text{Riemannian Hessian } \text{Hess}_y f[\xi_y] = \Pi_y(\text{D grad}_y f[\xi_y]),$$

where Π_y is the orthogonal projection operator (8).

D. Riemannian trust-region algorithm

The Riemannian trust-region (TR) algorithm approximates the function f with a *second-order* model at every iteration. The second-order model (which is called the trust-region sub-problem) makes use of the Riemannian gradient and Hessian computations as shown in Section IV-C. The trust-region sub-problem is then solved efficiently (using an iterative quadratic optimization solver, *e.g.*, with the truncated conjugate gradient algorithm) to obtain a candidate search direction. If the candidate search leads to an appreciable decrease in the function f , then it is accepted else it is rejected [23, Chapter 7]. Algorithm 1 summarizes the key steps of the proposed trust-region algorithm for solving (5).

Algorithm 1 Riemannian trust-region algorithm for (5)

Input: n -dimensional hyperbolic embeddings \mathbf{X} and rank r .
Initialize $y \in \mathcal{M} := \text{St}(r, n) \times \mathbb{H}^r \times \dots \times \mathbb{H}^r$.

repeat

1: Compute $\nabla_y f$.

2: **Riemannian TR step:** compute a search direction ξ_y which minimizes the trust region sub-problem. It makes use of $\nabla_y f$ and its directional derivative, and their Riemannian counterparts (10).

3: Update $y_+ = R_y(\xi_y)$ (retraction step) from (9).

until convergence

Output: $y = (\mathbf{U}, \bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_m)$ and $\hat{\mathbf{Z}} = \begin{bmatrix} 1 & \mathbf{0}_r^\top \\ \mathbf{0}_n & \mathbf{U} \end{bmatrix} \bar{\mathbf{Z}}$.

E. Computational complexity

The manifold-related ingredients cost $O(nr^2 + mr)$. For example, the computation of the Riemannian gradient in (10) involves only the tangent space projection operation that costs $O(nr^2 + mr)$. Similarly, the retraction operation costs $O(nr^2 + mr + r^3)$.

The computations of f and its derivatives cost $O(nmr)$ (for all the three choices of the loss function in Section III). The overall computational cost per iteration of our implementation is, therefore, $O(nmr)$.

F. Numerical implementation

We use the Matlab toolbox Manopt [25] to implement Algorithm 1 for (5). Manopt comes with a well-implemented generic Riemannian trust-region solver, which can be used appropriately to solve (5) by providing the necessary optimization-related ingredients mentioned earlier. The Matlab codes are available at <https://pratikjawanpuria.com>.

V. EXPERIMENTS

In this section, we evaluate the performance of the proposed low-rank parameterization of hyperbolic embeddings. In particular, we compare the quality of the low-rank hyperbolic embeddings obtained by minimizing the three different loss functions discussed in Section III.

Experimental setup and evaluation metric

We are provided with the hyperbolic embeddings corresponding to a hierarchical entity such as nodes of a tree or a graph. We also have the ground truth information of the given tree (or graph). Let $(\mathcal{G}, \mathcal{E})$ represents the ground truth, where \mathcal{G} is the set of nodes and $\mathcal{E} = \{(u, v)\}$ be the set of edges between the nodes ($u, v \in \mathcal{G}$). Hyperbolic embeddings can be employed to reconstruct the ground truth since a low hyperbolic distance (3) between a pair of nodes implies a high probability of an edge between them. However, such a reconstruction may also incorporate errors such as missing out on an edge or adding a non-existent edge.

We measure the quality of the hyperbolic embeddings as follows: let u and v be a pair of nodes in \mathcal{G} such that $(u, v) \in$

TABLE I

MEAN AVERAGE PRECISION (MAP) SCORE OBTAINED BY THE PROPOSED APPROACHES ON THE MAMMAL DATASET.

Rank	Method-1	Method-2	Method-3
5	0.8274	0.8416	0.1516
10	0.9106	0.9143	0.9004
20	0.9388	0.9390	0.9388
50	0.9488	0.9486	0.9488
100	0.9504	0.9504	0.9504
200	0.9502	0.9502	0.9502
300	0.9501	0.9501	0.9501

\mathcal{E} . Let \bar{u} and \bar{v} be the hyperbolic embeddings corresponding to u and v , respectively. We compute the hyperbolic distance $d_{\mathbb{H}}(\bar{u}, \bar{v})$ (3) and rank it among the distance corresponding to all untrue edges from u , i.e., $\{d_{\mathbb{H}}(u, w) | (u, w) \notin \mathcal{E}\}$. We then compute the mean average precision (MAP) of the ranking. The MAP score is a commonly employed metric for evaluating graph embeddings [3], [13], [26], [27]. Overall, we compare the quality of the proposed low-rank approximation by comparing the MAP score of the original high dimensional embeddings and the low-rank embeddings.

We obtain the original hyperbolic embeddings from the implementation provided by [3]. It should be noted that [3] learns the hyperbolic embeddings from the Poincaré model and we employ the transformation discussed in Section II to obtain embeddings corresponding to the hyperboloid model. It should be mentioned that though [13] directly learns hyperbolic embeddings from the hyperboloid model, its implementation is not available.

Datasets

We perform experiments on the mammal and noun subtrees of the WordNet database[20]. WordNet is a lexical database and among other things, it also provides relations between pairs of concepts.

The ‘mammal’ dataset has *mammal* as the root node, with ‘is-a’ (hyponymy) relationship defining the edges. As an example, it has relationships such as ‘rodent’ *is-a* ‘mammal’, ‘squirrel’ *is-a* ‘rodent’, *etc.* Hence, there exists an edge from the ‘mammal’ node to ‘rodent’ node and from ‘rodent’ node to ‘squirrel’ node. The WordNet mammal subtree consists of $|\mathcal{G}| = 1180$ nodes and $|\mathcal{E}| = 6540$ edges. A part of this subtree is displayed in Figure 1(a).

Similarly, the ‘noun’ dataset is also a subtree of WordNet database. Examples in this subtree include ‘photograph’ *is-a* ‘object’, ‘bronchitis’ *is-a* ‘disease’, ‘disease’ *is-a* ‘entity’, *etc.* It consists of $|\mathcal{G}| = 82115$ nodes and $|\mathcal{E}| = 743086$ edges.

Results

We compare the performance of the proposed low-rank approximation of hyperbolic embeddings with the three loss functions discussed in Section III. Table I reports the results on the mammal dataset with different values of rank $r = \{5, 10, 20, 50, 100, 200, 300\}$. The original 300-dimensional

TABLE II

MEAN AVERAGE PRECISION (MAP) SCORE OBTAINED BY THE PROPOSED APPROACHES ON THE NOUN DATASET.

Rank	Method-1	Method-2	Method-3
5	0.5343	0.5422	0.5343
10	0.6742	0.6796	0.6742
20	0.7425	0.7449	0.7425
50	0.7887	0.7891	0.7887
100	0.8070	0.8070	0.8070

hyperbolic embeddings for the mammal subtree achieve a MAP score of 0.9501. We observe that all the three methods are able to obtain MAP scores very close to the original embeddings with rank $r \geq 50$. In addition, Method-1 and Method-2 perform well even in very low-rank setting ($r = 5$). This hints that penalizing with the Euclidean distance may be more suitable than compared to the hyperbolic distance (3) for approximating hyperbolic embeddings when the given rank is very small.

The results on the noun dataset are reported in Table II. This dataset is challenging because of its scale and relatively low reconstruction performance of the original hyperbolic embeddings. The original 100-dimensional hyperbolic embeddings for the noun subtree achieve a MAP score of 0.8070. We observe that at rank $r = 20$ our methods are able to get within 90% of the performance obtained by the original embeddings.

VI. CONCLUSION AND FUTURE WORK

Recently, hyperbolic embeddings have gained popularity in many machine learning applications because of their ability to model complex networks. In this paper, we have looked at scenarios where hyperbolic embeddings are potentially high dimensional and how to compress them using a low-rank factorization model. While low-rank decomposition of Euclidean embeddings are well-known, that of hyperbolic embeddings has not been well-studied. To this end, we have proposed a systematic approach to compute low-rank approximations of hyperbolic embeddings. Our approach allows to decompose a high dimensional hyperbolic embedding (\bar{x}) into a product of low-dimensional subspace (\mathbf{U}) and a smaller dimensional hyperbolic embedding (\bar{z}).

We modeled the learning problem as an optimization problem on manifolds. Various optimization-related notions were presented to implement a Riemannian trust-region algorithm. Our experiments showed the benefit of the proposed low-rank approximations on real-world datasets.

As a future research direction, we would like to explore how low-rank hyperbolic embeddings are useful in downstream applications. Another research direction could be on developing methods to compute a “good” rank of hyperbolic embeddings.

REFERENCES

- [1] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [2] A. Muscoloni, J. M. Thomas, S. Ciucci, G. Bianconi, and C. V. Cannistraci. Machine learning meets complex networks via coalescent embedding in the hyperbolic space. *Nature Communications*, 8(1):1615, 2017.
- [3] M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Neural Information Processing Systems Conference (NIPS)*, 2017.
- [4] F. Sala, C. D. Sa, A. Gu, and C. Ré. Representation tradeoffs for hyperbolic embeddings. In *International Conference on Machine Learning (ICML)*, 2018.
- [5] A. Gu, F. Sala, B. Gunel, and C. Ré. Learning mixed-curvature representations in product spaces. In *International Conference on Learning Representations (ICLR)*, 2019.
- [6] B. Dhingra, C. J. Shallue, M. Norouzi, A. M. Dai, and G. E. Dahl. Embedding text in hyperbolic spaces. In *Twelfth Workshop on Graph-Based Methods for Natural Language Processing (ACL)*, pages 59–69, 2018.
- [7] A. Tifrea, G. Bécigneul, and O.-E. Ganea. Poincaré glove: Hyperbolic word embeddings. In *International Conference on Learning Representations (ICLR)*, 2019.
- [8] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.
- [9] M. Hamann. On the tree-likeness of hyperbolic spaces. *Mathematical Proceedings of the Cambridge Philosophical Society*, 164(2):345–361, 2018.
- [10] Y. Tay, L. A. Tuan, and S. C. Hui. Hyperbolic representation learning for fast and efficient neural question answering. In *Web Search and Data Mining (WSDM)*, 2018.
- [11] T. D. Q. Vinh, Y. Tay, S. Zhang, G. Cong, and X.-L. Li. Hyperbolic recommender systems. Technical report, arXiv:1809.01703, 2018.
- [12] B. P. Chamberlain, S. R. Hardwick, D. R. Wardrope, F. Dzongang, F. Daolio, and S. Vargas. Scalable hyperbolic recommender systems. Technical report, arXiv:1902.08648, 2019.
- [13] M. Nickel and D. Kiela. Learning continuous hierarchies in the Lorentz model of hyperbolic geometry. In *International Conference on Machine Learning (ICML)*, 2018.
- [14] O.-E. Ganea, G. Bécigneul, and T. Hofmann. Hyperbolic entailment cones for learning hierarchical embeddings. In *International Conference on Machine Learning (ICML)*, 2018.
- [15] O.-E. Ganea, G. Bécigneul, and T. Hofmann. Hyperbolic neural networks. In *Neural Information Processing Systems Conference (NIPS)*, 2018.
- [16] B. P. Chamberlain, J. R. Clough, and M. P. Deisenroth. Neural embeddings of graphs in hyperbolic space. Technical report, arXiv:1705.10359, 2017.
- [17] C. Gulcehre, M. Denil, M. Malinowski, A. Razavi, R. Pascanu, K. M. Hermann, P. Battaglia, V. Bapst, D. Raposo, A. Santoro, and N. Freitas. Hyperbolic attention networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [18] J. Anderson. *Hyperbolic Geometry*. Springer-Verlag London, 2005.
- [19] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Springer-Verlag New York, 2006.
- [20] G. A. Miller. Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.
- [21] B. Wilson and M. Leimeister. Gradient descent in hyperbolic space. Technical report, arXiv:1805.08207, 2018.
- [22] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
- [23] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.
- [24] X. Pennec. Hessian of the riemannian squared distance. Technical report, Université Côte d’Azur and Inria Sophia-Antipolis Méditerranée, 2017.
- [25] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15(Apr):1455–1459, 2014.
- [26] A. Bordes, N. Usunier, A. García-Durán, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *Neural Information Processing Systems Conference (NIPS)*, 2013.
- [27] M. Nickel, L. Rosasco, and T. A. Poggio. Holographic embeddings of knowledge graphs. In *AAAI Conference on Artificial Intelligence*, 2016.