

CO-SNE: Dimensionality Reduction and Visualization for Hyperbolic Data

Yunhui Guo

Haoran Guo
UC Berkeley / ICSI

Stella X. Yu

Abstract

Hyperbolic space can naturally embed hierarchies that often exist in real-world data and semantics. While high-dimensional hyperbolic embeddings lead to better representations, most hyperbolic models utilize low-dimensional embeddings, due to non-trivial optimization and visualization of high-dimensional hyperbolic data.

We propose CO-SNE, which extends the Euclidean space visualization tool, t-SNE, to hyperbolic space. Like t-SNE, it converts distances between data points to joint probabilities and tries to minimize the Kullback-Leibler divergence between the joint probabilities of high-dimensional data X and low-dimensional embedding Y . However, unlike Euclidean space, hyperbolic space is inhomogeneous: A volume could contain a lot more points at a location far from the origin. CO-SNE thus uses hyperbolic normal distributions for X and hyperbolic Cauchy instead of t-SNE’s Student’s t -distribution for Y , and it additionally seeks to preserve X ’s individual distances to the Origin in Y .

We apply CO-SNE to naturally hyperbolic data and supervisedly learned hyperbolic features. Our results demonstrate that CO-SNE deflates high-dimensional hyperbolic data into a low-dimensional space without losing their hyperbolic characteristics, significantly outperforming popular visualization tools such as PCA, t-SNE, UMAP, and HoroPCA which is also designed for hyperbolic data.

1. Introduction

Datasets with hierarchical structures are ubiquitous. Social networks [9] and complex networks [1] are representative examples of hierarchical data. Euclidean space cannot embed entities in such hierarchical datasets without distortion. Hyperbolic space, a non-Euclidean space with constant negative curvature, has been widely used for embedding hierarchical data since hyperbolic metric can closely approximate the tree metric. Hyperbolic space has thus been used in the representation learning of word embeddings [22] (Figure 1) and visual inputs [12, 17]. Several algorithms also directly operate on hyperbolic space [5, 6, 31].

Higher embedding dimension can generally lead to bet-

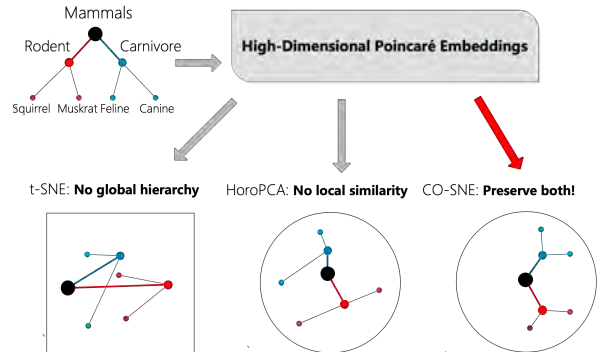


Figure 1. Our CO-SNE method deflates a high-dimensional hyperbolic representation while preserving their global hierarchy and local similarities. We generate the Poincaré embeddings [22] of the mammal subtree from WordNet in a five-dimensional hyperbolic space. We apply the standard t-SNE [30], the recently proposed HoroPCA [4] and our proposed CO-SNE method to visualize the embeddings in a two-dimensional Euclidean (in t-SNE) or hyperbolic space (in HoroPCA and CO-SNE). CO-SNE preserves the global hierarchy (root node is in the center and leaf nodes are close to boundary) and the local similarity (sibling nodes are close) in the two-dimensional embeddings.

ter hyperbolic representation quality [10, 22]. However, learning with two-dimensional hyperbolic space is prevalent [10, 22]. One of the reasons is the ease of visualization. There are several isometrically equivalent models for representing hyperbolic space. The Poincaré ball model is arguably the most widely used model in hyperbolic representation learning [8, 10, 22]. With the Poincaré ball model, we can easily visualize two-dimensional hyperbolic embeddings within a unit Euclidean circle. However, to visualize high-dimensional hyperbolic data is not easy, as most visualization methods assume the data exist in Euclidean space.

Embeddings in Poincaré ball have two notable properties: 1) The embeddings have a global hierarchical structure. Root nodes are in the center of the ball and leaf nodes are close to the boundary of the ball. 2) The embeddings possess a local similarity structure. Sibling nodes should be close in the embedding space.

t-SNE [30] is a popular visualization tool for visualiz-

ing high-dimensional Euclidean data. However, t-SNE cannot preserve the global hierarchy of the hyperbolic embeddings. HoroPCA [4] is recently proposed as an extension of PCA on hyperbolic space. However, HoroPCA cannot preserve the local similarity of the hyperbolic embeddings. In this paper, we propose CO-SNE which can preserve both the global hierarchy and local similarity of high-dimensional hyperbolic embeddings in a low-dimensional hyperbolic space (1). In Figure 1, the *mammal* subtree from WordNet [20] is embedded in a five-dimensional hyperbolic space via Poincaré embeddings [22]. We use t-SNE [30], HoroPCA [4] and CO-SNE to visualize the embeddings in a two-dimensional space. CO-SNE can preserve the structure of the data well.

In CO-SNE, for maintaining local similarity structure, we adopt the same idea as in the standard t-SNE to minimize the KL-divergence between the high-dimensional similarities and the low-dimensional similarities. We adopt hyperbolic versions of the normal distribution and Cauchy distribution to compute the similarities. To maintain the global hierarchical structure, we adopt a distance loss function which seeks to preserve the individual distances of high-dimensional hyperbolic embeddings to the Origin in low-dimensional hyperbolic space.

In summary, we make the following contributions,

- We propose CO-SNE which can represent high-dimensional hyperbolic datapoints in a low-dimensional hyperbolic space while maintaining the local similarity and the global hierarchical structure.
- We propose to use hyperbolic Cauchy distribution for computing low-dimensional similarities which is crucial for producing good visualization in hyperbolic space.
- We apply CO-SNE to visualize synthetic hyperbolic data, hierarchical biological datasets and hyperbolic embeddings learned by supervised and unsupervised learning methods to better understand high-dimensional hyperbolic data. Across all the cases, CO-SNE produces much better visualization than the baselines.

2. Related Work

Large-margin Classification [5] in hyperbolic space was proposed by changing the distance function from Euclidean to hyperbolic. Robust large-margin classification [31] in hyperbolic space is also proposed and has the first theoretical guarantees for learning a hyperbolic classifier.

Hyperbolic Neural Networks (HNNs) [6] rewrite multinomial logistic regression (MLR), fully connected layers and Recurrent Neural Networks for hyperbolic embed-

dings via gyrovectors space operations [29]. In a follow-up work, Hyperbolic Neural Networks++ [27] introduces hyperbolic convolutional layers. Hyperbolic attention networks [7] are proposed by extending attention operations to hyperbolic space in a manner similar to [6]. Hyperbolic graph neural networks [16] are further proposed by altering the geometry of Graph Neural Networks (GNNs) [33] to hyperbolic space. Hyperbolic neural networks have also been used for visual inputs and achieve better results than Euclidean neural networks on tasks such as few-shot classification and person re-identification [12].

Unsupervised Learning Methods in Hyperbolic Space are also proposed recently. Poincaré embeddings are proposed to embed words and relations in hyperbolic space [22]. [21] proposed a generalized version of the normal distribution on hyperbolic space called wrapped normal distribution. The proposed wrapped normal distribution is used as the latent space for constructing hyperbolic variational autoencoders (VAEs) [13]. A similar idea is also adopted in [17] to construct Poincaré VAEs. Unsupervised 3D segmentation [10] and instance segmentation [32] are achieved in hyperbolic space via hierarchical triplet loss.

Data Visualization is the process of generating low-dimensional representation of each high-dimensional datapoint. A good visualization should maintain the interesting structure of the data presented in high-dimensional space. t-distributed stochastic neighbor embedding (t-SNE) [30] is arguably the most widely used tool for data visualization. t-SNE attempts to maintain the local similarities of the high-dimensional datapoints in the low-dimensional space. More recently, UMAP [18] is proposed as a manifold learning technique for dimensionality reduction and visualization. Compared with t-SNE, UMAP can better preserve the global structure of the high-dimensional data.

Dimensionality Reduction methods also can be used for data visualization. Principal Component Analysis (PCA [11] is a commonly used dimensionality reduction technique. PCA aims at maintaining the maximum amount of variation of information in the low-dimensional space. Isomap [28] is a non-linear dimensionality reduction method which attempts to preserve local structures. LLE [26] is another popular non-linear dimensionality reduction method that can produce neighborhood preserving embeddings of high-dimensional data.

Notably, no existing data visualization methods are designed for visualizing high-dimensional hyperbolic data. Our proposed CO-SNE method can be used to visualize high-dimensional hyperbolic embeddings.

3. CO-SNE

We present the proposed CO-SNE method which can faithfully represent high-dimensional hyperbolic data in a low-dimensional hyperbolic space. In Section 3.1, we re-

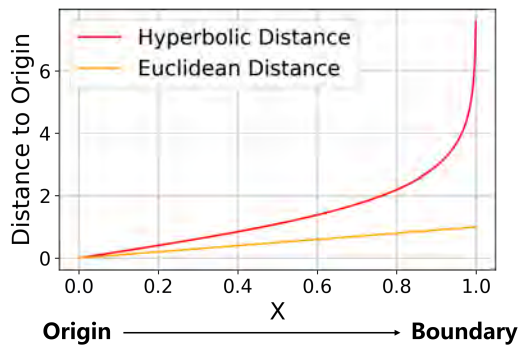


Figure 2. The comparison of Euclidean distance and hyperbolic distance as we move towards the boundary of the Poincaré ball.

view the basics of Poincaré ball model for hyperbolic space. In Section 3.2, we review t-distributed Stochastic neighbor embedding (t-SNE). In Section 3.3 and 3.4, we introduce the hyperbolic version of normal distribution and Student’s t-distribution. In Section 3.5, we discuss the issue of using hyperbolic Student’s t-distribution for computing low-dimensional similarities and propose to use hyperbolic Cauchy distribution. In Section 3.6, we present the distance loss for maintaining global hierarchy. In Section 3.7, we present the optimization details of CO-SNE.

3.1. Poincaré Ball Model for Hyperbolic Space

A hyperbolic space is a Riemannian manifold with a constant negative curvature. There are several isometrically equivalent models for representing hyperbolic space. The Poincaré ball model is the most commonly used one in hyperbolic representation learning [6, 22]. The n -dimensional Poincaré ball model is defined as $(\mathbb{B}^n, \mathfrak{g}_{\mathbf{x}})$, where $\mathbb{B}^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1\}$ and $\mathfrak{g}_{\mathbf{x}} = (\gamma_{\mathbf{x}})^2 I_n$ is the Riemannian metric tensor. $\gamma_{\mathbf{x}} = \frac{2}{1 - \|\mathbf{x}\|^2}$ is the conformal factor and I_n is the Euclidean metric tensor. Given two points $\mathbf{u} \in \mathbb{B}^n$ and $\mathbf{v} \in \mathbb{B}^n$, the hyperbolic distance between them is defined as,

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

where arcosh is the inverse hyperbolic cosine function and $\|\cdot\|$ is the usual Euclidean norm. Different from Euclidean distance, hyperbolic distance grows exponentially fast as we move the points towards the boundary of the Poincaré ball as in Figure 2.

3.2. t-SNE

t-SNE [30] begins by mapping high-dimensional distances between datapoints to similarity values. The similarity values are either conditional or joint probabilities based on the probability a point will pick another as its neighbor if

neighbors are picked in proportion to the probability density of a distribution centered at that point. t-SNE defines the conditional probability $p_{j|i}$, the probability that the point \mathbf{x}_i will pick a point \mathbf{x}_j as its neighbor, using a normal distribution centered at the point \mathbf{x}_i . t-SNE then defines the joint probability distribution P , by setting $p_{ij} = \frac{p_{i|j} + p_{j|i}}{2m}$ as a way to increase the cost contribution of outlier points, where m is the number of high-dimensional datapoints. The conditional probability density $p_{j|i}$ is,

$$p_{j|i} = \frac{\exp(-d(\mathbf{x}_i, \mathbf{x}_j)^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-d(\mathbf{x}_i, \mathbf{x}_k)^2 / 2\sigma_i^2)} \quad (2)$$

where $d(\mathbf{x}_i, \mathbf{x}_j)$ is the distance between \mathbf{x}_i and \mathbf{x}_j . In the low-dimensional space, Student’s t-distribution is used for modeling the joint probability distribution Q between embeddings, and q_{ij} is defined as,

$$q_{ij} = \frac{(1 + d(\mathbf{y}_i, \mathbf{y}_j)^2)^{-1}}{\sum_{k \neq l} (1 + d(\mathbf{y}_k, \mathbf{y}_l)^2)^{-1}} \quad (3)$$

where \mathbf{y}_i is the corresponding low-dimensional embedding of \mathbf{x}_i . In t-SNE, to maintain the local similarities, the cost function to minimize is the Kullback-Leibler divergence between the probability distributions P and Q :

$$\mathcal{C} = KL(P||Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \quad (4)$$

One direct extension of t-SNE is to replace Euclidean version of normal distribution and Student’s t-distribution with hyperbolic normal distribution and hyperbolic Student’s t-distribution. We call such a direct extension as HT-SNE. In Section 3.3 and Section 3.4, we show how to generalize the distributions to hyperbolic space.

3.3. Hyperbolic Normal Distribution

To define the conditional probability in the high-dimensional hyperbolic space, we need to generalize the normal distribution to hyperbolic space. One natural generalization is called the Riemannian normal distribution which is the maximum entropy probability distribution given an expectation and a variance [25]. Given the Fréchet mean $\boldsymbol{\mu} \in \mathbb{B}_c^n$ and a dispersion parameter $\sigma > 0$, the Riemannian normal distribution is defined as,

$$\mathcal{N}_{\mathbb{B}^n}(\mathbf{x}|\boldsymbol{\mu}, \sigma^2) = \frac{1}{\mathbf{Z}} \exp\left(-\frac{d_{\mathbb{B}^n}(\boldsymbol{\mu}, \mathbf{x})^2}{2\sigma^2}\right) \quad (5)$$

where \mathbf{Z} is the normalization constant. There are other generalizations of the normal distribution in hyperbolic space [17]. We use the Riemannian normal distribution for simplicity. Thereafter, we refer to Riemannian normal distribution as hyperbolic normal distribution.

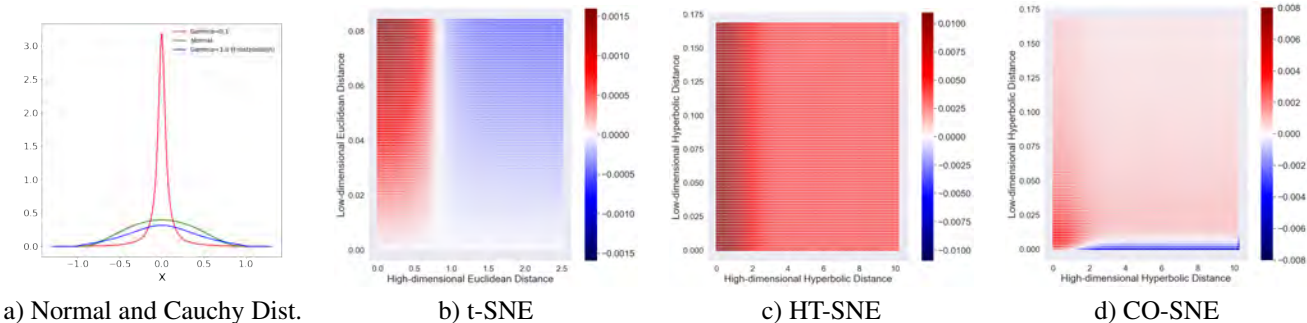


Figure 3. Hyperbolic Student’s t-distribution is not heavy-tailed which leads to strong attraction forces in HT-SNE. a) The probability density function of the hyperbolic normal distribution, hyperbolic Cauchy distribution with $\gamma = 0.1$ and hyperbolic Cauchy distribution with $\gamma = 1.0$. Hyperbolic Student’s t-distribution (hyperbolic Cauchy distribution with $\gamma = 1.0$) is not heavy-tailed. b) The gradients of the standard t-SNE as a function of low-dimensional and high-dimensional Euclidean distance. c) The gradients of the HT-SNE as a function of low-dimensional and high-dimensional hyperbolic distance. d) The gradients of CO-SNE as a function of low-dimensional and high-dimensional hyperbolic distance. There is strong repulsion when dissimilar high-dimensional datapoints are projected close.

3.4. Hyperbolic Student’s t-Distribution

One way to define the Student’s t-distribution is to express the random variable t as,

$$t = \frac{u}{\sqrt{v/n}} \quad (6)$$

where u is a random variable sampled from a standard normal distribution and v is a random variable sampled from a χ^2 -distribution of n degrees of freedom. In particular, t-SNE adopts a Student’s t-distribution with one degree of freedom and the probability density function is defined as,

$$f(t; t_0) = \frac{1}{\pi(1 + (t - t_0)^2)} \quad (7)$$

To extend the Student’s t-distribution to hyperbolic space, we derive the probability density function as,

$$f_{\mathbb{B}^n}(t; t_0) = \frac{1}{\pi(1 + d_{\mathbb{B}^n}(t, t_0)^2)} \quad (8)$$

The details can be found in the Supplementary.

3.5. Hyperbolic t-Distribution is Not Enough

The motivation to use Student’s t-distribution in the standard t-SNE is that Student’s t-distribution has heavier tails than the normal distribution. This causes a repulsion force between dissimilar high-dimensional datapoints which helps mitigate the “Crowding Problem” [30]. However, hyperbolic Student’s t-distribution is not heavy-tailed since the hyperbolic distance grows exponentially fast.

Figure 3 plots the hyperbolic normal distribution and hyperbolic Student’s t-distribution. It can be noted that the tails of hyperbolic Student’s t-distribution diminish as fast as hyperbolic normal distribution. In contrast, the standard Euclidean Student’s t-distribution has heavier tails than the

normal distribution. The consequence is that the repulsion force and the attraction force with hyperbolic Student’s t-distribution behave drastically different from that of using Euclidean Student’s t-distribution.

Similar to [30], we plot the gradients between two low-dimensional embeddings \mathbf{y}_i and \mathbf{y}_j as a function of their pairwise Euclidean (hyperbolic) distance in the low-dimensional Euclidean (hyperbolic) and the pairwise Euclidean (hyperbolic) distance of the corresponding datapoints in the high-dimensional Euclidean (hyperbolic) space. Positive values indicate attraction and negative values indicate repulsion between the embeddings \mathbf{y}_i and \mathbf{y}_j respectively. Here we only consider the initial stage when all the low-dimensional embeddings are close. The results are shown in Figure 3. We have two observations. First, for the standard t-SNE, there is strong repulsion when dissimilar high-dimensional datapoints are projected closely. Second, with the hyperbolic Student’s t-distribution in HT-SNE, there is little or no repulsion force which causes the low-dimensional embeddings tend to cluster together.

To remedy this issue, we seek to replace the hyperbolic Student’s t-distribution to create more repulsion forces between dissimilar high-dimensional datapoints. Consider the hyperbolic Cauchy distribution which has the probability density function,

$$f(t; t_0, \gamma) = \frac{1}{\pi\gamma} \left[\frac{\gamma^2}{d_{\mathbb{B}^n}(t, t_0)^2 + \gamma^2} \right] \quad (9)$$

where γ is the scale parameter. Notice that the Student’s t-distribution is a special case of Cauchy distribution with $\gamma = 1.0$. With a small γ , Cauchy distribution has a higher peak (Figure 3 (a)). Thus, if the high-dimensional datapoints are modeled with close low-dimensional embeddings, the density value is much larger than the corresponding density in the high-dimensional space. This produces

	Metric	High/low-dimensional Dist.	Losses
t-SNE	Euclidean	Normal/t-distribution	KL-div
HT-SNE	Hyperbolic	Normal/t-distribution	KL-div
CO-SNE	Hyperbolic	Normal/Cauchy	KL-div + Distance

Table 1. Our CO-SNE extends t-SNE by adopting hyperbolic normal distribution and hyperbolic Cauchy distribution. Compared with t-SNE, CO-SNE assumes the high-dimensional and low-dimensional space are hyperbolic. t-SNE cannot maintain the global hierarchy of the hyperbolic embeddings. Compared with HT-SNE, CO-SNE adopts hyperbolic Cauchy distribution and an additional distance loss. HT-SNE cannot push dissimilar high-dimensional points away in the low-dimensional space.

a strong repulsion when dissimilar high-dimensional datapoints are modeled with close low-dimensional embeddings. More details can be found in the Supplementary. We show the corresponding gradients in Figure 3 (d). We can observe that there is a strong repulsion force when high-dimensional dissimilar datapoints are projected closely in CO-SNE. Such repulsion is particularly important in the initial stage of training since all the low-dimensional embeddings are close.

3.6. Distance Loss

Hyperbolic space can be naturally used for embedding tree structured data. The datapoint which is close to the center of the Poincaré ball can be viewed as the root node and datapoints which are close to the boundary of the Poincaré ball can be viewed as leaf nodes. The criterion of t-SNE cannot preserve the tree-structure of the high-dimensional datapoints in the low-dimensional space. One natural way to characterize the level of the datapoint in the underlying tree is the norm of the datapoint. Thus, we attempt to keep the norm invariant after the high-dimensional datapoint is projected. This leads to the following loss function which minimizes the difference between the norms of the high-dimensional datapoint \mathbf{x}_i and the corresponding low-dimensional embedding \mathbf{y}_i ,

$$\mathcal{H} = \frac{1}{m} \sum_{i=1}^m (\|\mathbf{x}_i\|^2 - \|\mathbf{y}_i\|^2)^2 \quad (10)$$

With the distance loss, we can preserve the global hierarchy of the high-dimensional hyperbolic embeddings.

3.7. Optimization of CO-SNE

Criterion. The criterion of CO-SNE is composed of the KL-divergence for maintaining local similarity and the distance loss for maintaining the global hierarchy,

$$\mathcal{L} = \lambda_1 \mathcal{C} + \lambda_2 \mathcal{H} \quad (11)$$

where λ_1 and λ_2 are hyperparameters. Table 1 shows a comparison of t-SNE, HT-SNE and CO-SNE.

Gradients. The cost function is optimized by gradient descent with respect to the low-dimensional embeddings. The gradient of the KL-divergence \mathcal{C} with respect to \mathbf{y}_i is given by,

$$\begin{aligned} \frac{\delta \mathcal{C}}{\delta \mathbf{y}_i} &= \sum_j \frac{\delta \mathcal{C}}{\delta d_{\mathbb{B}^n}(\mathbf{y}_i, \mathbf{y}_j)} \frac{\delta d_{\mathbb{B}^n}(\mathbf{y}_i, \mathbf{y}_j)}{\delta \mathbf{y}_i} \\ &= 2 \sum_j (p_{ij} - q_{ij}) (1 + d_{\mathbb{B}^n}(\mathbf{y}_i, \mathbf{y}_j)^2)^{-1} \frac{\delta d_{\mathbb{B}^n}(\mathbf{y}_i, \mathbf{y}_j)}{\delta \mathbf{y}_i} \end{aligned} \quad (12)$$

The partial gradient of the distance with respect to the low dimensional embedding \mathbf{y}_i is given by

$$\frac{\delta d_{\mathbb{B}^n}(\mathbf{y}_i, \mathbf{y}_j)}{\delta \mathbf{y}_i} = \frac{4}{\beta \sqrt{\gamma^2 - 1}} \left(\frac{\|\mathbf{y}_j\|^2 - 2\langle \mathbf{y}_i, \mathbf{y}_j \rangle + 1}{\alpha^2} \mathbf{y}_i - \frac{\mathbf{y}_j}{\alpha} \right) \quad (13)$$

where $\alpha = 1 - \|\mathbf{y}_i\|^2$, $\beta = 1 - \|\mathbf{y}_j\|^2$, $\gamma = 1 + \frac{2}{\alpha\beta} \|\mathbf{y}_i - \mathbf{y}_j\|^2$. We use the Riemannian stochastic gradient descent [3] for the KL-divergence \mathcal{C} . The gradient of the distance loss \mathcal{H} with respect to \mathbf{y}_i is computed as,

$$\frac{\delta \mathcal{H}}{\delta \mathbf{y}_i} = -4(\|\mathbf{x}_i\|^2 - \|\mathbf{y}_i\|^2) \mathbf{y}_i \quad (14)$$

We constrain the embeddings to the Poincaré ball after each update as in [22].

Stagewise Training. We found the following training strategy is useful for producing better visualization in practice. We train the low-dimensional embeddings only with the local similarity loss for the first 500 iterations. Then we add the distance loss. The stopping criterion is the same as in t-SNE [30]. The reason of the strategy is that it is hard to move the points when they are approaching the boundary of the Poincaré ball.

4. Experiments and Results

Baselines. We compare CO-SNE with 4 baselines,

- The standard t-SNE [30]: this is the standard t-SNE which adopts Euclidean distance for computing the similarities in the high-dimensional space and the low-dimensional space.
- Principal Component Analysis (PCA) [11]: PCA is a commonly used dimensionality reduction method which attempts to maintain the maximal variation of the data. However, as a linear dimensionality method, PCA cannot reduce high-dimensional data to two dimensions in a meaningful way [2].
- HoroPCA [4]: HoroPCA is a recently proposed extension of PCA on hyperbolic space. HoroPCA proposed to parameterize geodesic subspaces by ideal points in

the Poincaré ball. HoroPCA also generalized the notion of projection and the objective function of PCA to hyperbolic space. HoroPCA can be used as a data whitening method for hyperbolic data and as a visualization method as well.

- UMAP [18]: UMAP is a recently proposed dimensionality reduction and visualization method based on the ideas from Riemannian geometry and topology. UMAP is competitive with t-SNE and can better preserve the global structure of the data.

Tasks. We consider 4 datasets to validate the effectiveness of CO-SNE: 1) synthetic datasets sampled from a mixture of hyperbolic normal distributions, 2) cellular differentiation data, 3) supervised hyperbolic embeddings and 4) unsupervised hyperbolic embeddings.

Implementation Details. For PCA and the standard t-SNE, we use the implementation from [24]. For UMAP, we use the implementation from [19]. For HoroPCA, we use the implementation from <https://github.com/HazyResearch/HoroPCA> which is provided by the original authors. We implement CO-SNE based on t-SNE by modifying the way of computing similarities and the optimization procedure.

Hyperparameters. For PCA, UMAP and HoroPCA, we use the default hyperparameters. For SNE-based methods, we initialize the low-dimensional embedding using a normal distribution with mean 0.01 and unit variance. The training follows the standard setups as in [24]. For CO-SNE, the scaling parameter γ is set to 0.1. The hyperparameter λ_1 is usually set to 10.0 and the hyperparameter λ_2 is usually set to 0.01.

4.1. Synthetic Point Clusters

We first use a synthetic dataset to validate the efficacy of CO-SNE. We randomly generate five point clusters of 20 points each in the 5D hyperbolic space. Each follows a hyperbolic normal distribution with the unit variance and the mean located on a different axis. The first and second means are close to the origin, at $[0.1, 0, 0, 0, 0]$ and $[0, -0.2, 0, 0, 0]$ respectively, whereas the third and fourth means are far from and equi-distance to the origin, at $[0, 0, 0.9, 0, 0]$ and $[0, 0, 0, -0.9, 0]$ respectively. The last mean is right at the origin $[0, 0, 0, 0, 0]$. Figure 4 shows the 2D visualization results of these 5D points by different methods.

CO-SNE produces a much better visualization compared with the baselines. With CO-SNE, the projected two-dimensional hyperbolic embeddings can preserve the local similarity structure and the global hierarchical structure of high-dimensional datapoints well. Also, CO-SNE can prevent the high-dimensional datapoints from being projected too close which usually happens with Euclidean distance based methods. Note that HT-SNE does not have





Figure 4. The projection of high-dimensional hyperbolic datapoints sampled from a mixture of hyperbolic normal distributions in a two-dimensional space with different methods. CO-SNE produces two-dimensional hyperbolic embeddings which preserve the hierarchical and similarity structure of the high-dimensional hyperbolic datapoints.


not enough repulsion to push close low-dimensional embeddings apart. We provide a detailed analysis on the drawbacks of each baseline for visualizing high-dimensional hyperbolic datapoints.

- Standard t-SNE: Euclidean distances are used for computing similarities in the high-dimensional space. Hyperbolic distances grow much faster than Euclidean distances. For high-dimensional hyperbolic datapoints which are close to boundary of the Poincaré ball, the standard t-SNE wrongly underestimates the distance between them. As a consequence, the standard t-SNE would take dissimilar high-dimensional data points as neighbors. The resulting low-dimensional embeddings would collapse into one point which leads to poor visualization. In summary, t-SNE cannot preserve the global hierarchy of the hyperbolic data.
- PCA and HoroPCA: as mentioned above, PCA and HoroPCA are linear dimensionality reduction methods which are not generally suitable for visualization in a two-dimensional space. Both PCA and HoroPCA cannot preserve local similarity of the hyperbolic data.
- UMAP: UMAP suffers from the same issue as t-SNE since Euclidean distance is used for computing high-dimensional similarities.

We mainly compare CO-SNE with HoroPCA since HoroPCA is specifically designed for hyperbolic data.



a) Canonical hematopoietic cell lineage tree [14]



b) HoroPCA

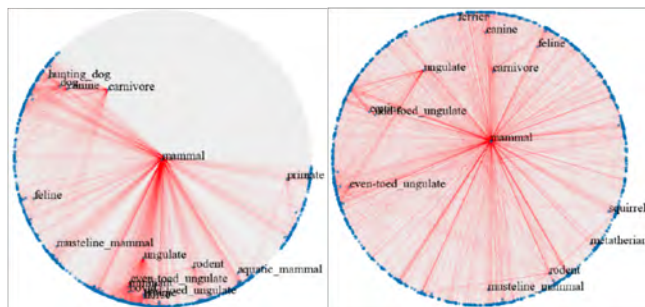
c) CO-SNE

Figure 5. CO-SNE produces better visualization of the high-dimensional biological datapoints than HoroPCA which captures the hierarchical and similarity structure in the data. a) The hierarchy of the original data from [14]. b) The two-dimensional embeddings generated by HoroPCA. c) The two-dimensional embeddings generated by CO-SNE.

4.2. Biological Dataset

Biological data can reveal naturally occurring hierarchies, such as in single cell RNA sequencing data [14]. In [14], they analyze cellular differentiation data, the transition of immature cells to specialized types. The immature cells can be viewed as the root of the tree and can branch off into several different types of cells, creating hierarchical data of cells in different states of progress in the transition process. One dataset we adapt from [14] is the mouse myelopoiesis dataset presented by [23], where there are 532 cells of 9 types. Two of the types, HSPC-1 and HSPC-2, form the root of the hierarchy, while megakaryocytic (Meg), erythrocytic (Eryth), monocyte-dendritic cell precursor (MDP), monocytic (Mono), and myelocyte (myelocytes and metamyelocytes) type cells are states father from the root. Granulocytic (Gran) cells are a precursor to myelocytes and multi-lineage primed (Multi-Lin) cells are in an intermediate state.

The data has originally 382 dimensions (noisy) and like [14] we first reduce it to 20 dimensions via PCA to reduce the noises. We then scale the data to fit within the Poincaré ball and run CO-SNE and HoroPCA to produce two-dimensional hyperbolic embeddings. Noted that this dataset is not centered. The results are shown in Figure 5.



a) HoroPCA

b) CO-SNE

Figure 6. CO-SNE produces better visualization of high-dimensional Poincaré word embeddings than HoroPCA. a) The two-dimensional Poincaré word embeddings generated from HoroPCA. b) The two-dimensional Poincaré word embeddings generated from CO-SNE. In the embeddings generated by CO-SNE, the word *feline* and *canine* are close to *canivore*, which is not the case in HoroPCA.

The proposed low-embeddings by CO-SNE capture the hierarchical structure in the original data.

4.3. Hierarchical Word Embeddings

Hyperbolic space has been used to embed hierarchical representations of symbolic data. In [22], the authors adopt hyperbolic space for embedding taxonomies, in particular, the transitive closure of WordNet noun hierarchy [20]. As shown in [22], higher-dimensional hyperbolic embeddings often lead to better representations, but they are harder to visualize. Following [22], we embed the hypernymy relations of the mammals subtree of WordNet in hyperbolic space. We use the open source implementation provided by [22] to train the ten-dimensional embeddings. We use HoroPCA and CO-SNE to visualize the learned embeddings in a two-dimensional hyperbolic space.

Figure 6 shows that compared with HoroPCA, CO-SNE can better preserve the hierarchical and similarity structure of the high-dimensional datapoints. For example, the word *feline* and *canine* are close to *canivore* in the CO-SNE embeddings, which is not the case in HoroPCA. The embeddings produced by CO-SNE also more resemble the two-dimensional embeddings as shown in Figure 2b of [22].

4.4. Features of Hyperbolic Neural Networks

We apply CO-SNE to visualize the embeddings produced by hyperbolic neural networks for supervised image classification. We train a hyperbolic neural network (HNN) [6] with feature clipping [8] on MNIST. The clipping value is 1.0 and the feature dimension is 64. We use HoroPCA and CO-SNE to reduce the dimensionality of the features to two. The full test set of MNIST cannot be used




Figure 7. CO-SNE produces better visualization of hyperbolic neural networks’ (HNNs) features than HoroPCA. a) The visualization produced by HoroPCA. b) The visualization produced by CO-SNE. In CO-SNE, the classes are well separated and have a clear hierarchical structure.




Figure 8. CO-SNE produces better visualization of the high-dimensional latent space representations generated by the encoder of the Poincaré VAE than HoroPCA. a) The two-dimensional latent Poincaré embeddings generated from HoroPCA. b) The two-dimensional latent Poincaré embeddings generated from CO-SNE. CO-SNE can capture the hierarchical and clustering structures of the high-dimensional latent representations while HoroPCA cannot.

due to the out-of-memory issue of HoroPCA. So for each class, we randomly sample 100 images.

Figure 7 shows the two-dimensional embeddings generated by HoroPCA and CO-SNE. The visualization produced by CO-SNE is significantly better the visualization produced by HoroPCA. In CO-SNE, the classes are well separated and have a clear hierarchical structure. We further train hyperbolic classifiers [6] on the frozen two-dimensional features. We use a learning rate of 0.001 and the number of epoch is 10. The accuracy of the features generated by HoroPCA is 30.2% while the accuracy of the features generated by CO-SNE is 61.2%. This implies that low-dimensional features generated by CO-SNE are more separated and respect the structure of original high-dimensional embeddings.

5. Latent Space of Poincaré VAE

Variational autoencoders (VAEs) [13] is a popular unsupervised learning method. Standard VAEs assume the latent space is Euclidean. [17] extends VAEs by assuming the latent space is hyperbolic. Different from the standard VAEs, Poincaré VAEs can embed tree-like structures more efficiently. For using CO-SNE to visualize the latent space of Poincaré VAEs, we train a Poincaré VAE with a latent dimension of five on MNIST [15] following from [17]. We further generate the latent space representations of 1000 randomly sampled images with the encoder.

Figure 8 shows the visualizations produced by HoroPCA and CO-SNE. Clearly, CO-SNE produces much better visualization than HoroPCA. In particular, we can easily observe the hierarchical and clustering structures in the latent space which are totally distorted in the visualization produced by HoroPCA. Thus, CO-SNE can be used to understand the latent space of Poincaré VAEs and facilitate the development of better unsupervised hyperbolic learning methods.

5.1. Discussion

CO-SNE can be used to visualize hyperbolic datapoints in a two-dimensional hyperbolic space while maintaining the local similarities and hierarchical structure. CO-SNE still suffers from the same weaknesses as t-SNE. In particular, as a general method for dimensionality reduction, the curse of intrinsic dimensionality and the non-convexity, please see [30] for more details. More ablations and results can be found in the Supplementary. In particular, we investigate the effect of hyperparameters for balancing the KL-divergence and the distance loss.

6. Summary

We propose CO-SNE which can be used to visualize hyperbolic data in a two-dimensional hyperbolic space. We use the generalized versions of the normal and Cauchy distributions in hyperbolic space to compute the high-dimensional and the low-dimensional similarities, respectively. We apply CO-SNE to visualize synthetic hyperbolic data, biological data and embeddings learned by supervised and unsupervised representation learning methods. CO-SNE produces much better visualizations than several popular visualization methods.

Acknowledgements. This work was supported, in part, by Berkeley Deep Drive and the National Science Foundation (NSF) under Grant No. 2131111. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

- [1] Gregorio Alanis-Lobato, Pablo Mier, and Miguel A Andrade-Navarro. Efficient embedding of complex networks to hyperbolic space via their laplacian. *Scientific reports*, 6(1):1–10, 2016. **1**
- [2] Sanjeev Arora, Wei Hu, and Pravesh K Kothari. An analysis of the t-sne algorithm for data visualization. In *Conference On Learning Theory*, pages 1455–1462. PMLR, 2018. **5**
- [3] Silvere Bonnabel. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013. **5**
- [4] Ines Chami, Albert Gu, Dat P Nguyen, and Christopher Ré. Horopca: Hyperbolic dimensionality reduction via horospherical projections. In *International Conference on Machine Learning*, pages 1419–1429. PMLR, 2021. **1, 2, 5**
- [5] Hyunghoon Cho, Benjamin DeMeo, Jian Peng, and Bonnie Berger. Large-margin classification in hyperbolic space. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1832–1840. PMLR, 2019. **1, 2**
- [6] Octavian-Eugen Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *arXiv preprint arXiv:1805.09112*, 2018. **1, 2, 3, 7, 8**
- [7] Caglar Gulcehre, Misha Denil, Mateusz Malinowski, Ali Razavi, Razvan Pascanu, Karl Moritz Hermann, Peter Battaglia, Victor Bapst, David Raposo, Adam Santoro, et al. Hyperbolic attention networks. *arXiv preprint arXiv:1805.09786*, 2018. **2**
- [8] Yunhui Guo, Xudong Wang, Yubei Chen, and Stella X Yu. Free hyperbolic neural networks with limited radii. *arXiv preprint arXiv:2107.11472*, 2021. **1, 7**
- [9] Mangesh Gupte, Pravin Shankar, Jing Li, Shanmugaelayut Muthukrishnan, and Liviu Iftode. Finding hierarchy in directed online social networks. In *Proceedings of the 20th international conference on World wide web*, pages 557–566, 2011. **1**
- [10] Joy Hsu, Jeffrey Gu, Gong-Her Wu, Wah Chiu, and Serena Yeung. Learning hyperbolic representations for unsupervised 3d segmentation. *arXiv preprint arXiv:2012.01644*, 2020. **1, 2**
- [11] Ian T Jolliffe and Jorge Cadima. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150202, 2016. **2, 5**
- [12] Valentin Khruikov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6418–6428, 2020. **1, 2**
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. **2, 8**
- [14] Anna Klimovskaia, David Lopez-Paz, Léon Bottou, and Maximilian Nickel. Poincaré maps for analyzing complex hierarchies in single-cell data. *Nature communications*, 11(1):1–9, 2020. **7**
- [15] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>. **8**
- [16] Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *arXiv preprint arXiv:1910.12892*, 2019. **2**
- [17] Emile Mathieu, Charline Le Lan, Chris J Maddison, Ryota Tomioka, and Yee Whye Teh. Continuous hierarchical representations with poincaré variational auto-encoders. *arXiv preprint arXiv:1901.06033*, 2019. **1, 2, 3, 8**
- [18] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. **2, 6**
- [19] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29):861, 2018. **6**
- [20] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. **2, 7**
- [21] Yoshihiro Nagano, Shoichiro Yamaguchi, Yasuhiro Fujita, and Masanori Koyama. A wrapped normal distribution on hyperbolic space for gradient-based learning. In *International Conference on Machine Learning*, pages 4693–4702. PMLR, 2019. **2**
- [22] Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *arXiv preprint arXiv:1705.08039*, 2017. **1, 2, 3, 5, 7**
- [23] Andre Olsson, Meenakshi Venkatasubramanian, Viren K Chaudhri, Bruce J Aronow, Nathan Salomonis, Harinder I Singh, and H. Leighton Grimes. Single-cell analysis of mixed-lineage states leading to a binary cell fate choice. *Nature*, 2016. **7**
- [24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. **6**
- [25] Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006. **3**
- [26] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000. **2**
- [27] Ryohei Shimizu, Yusuke Mukuta, and Tatsuya Harada. Hyperbolic neural networks++. *arXiv preprint arXiv:2006.08210*, 2020. **2**
- [28] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000. **2**
- [29] Abraham Albert Ungar. A gyrovector space approach to hyperbolic geometry. *Synthesis Lectures on Mathematics and Statistics*, 1(1):1–194, 2008. **2**
- [30] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. **1, 2, 3, 4, 5, 8**
- [31] Melanie Weber, Manzil Zaheer, Ankit Singh Rawat, Aditya Menon, and Sanjiv Kumar. Robust large-margin learning in hyperbolic space. *arXiv preprint arXiv:2004.05465*, 2020. **1, 2**

- [32] Zhenzhen Weng, Mehmet Giray Ogut, Shai Limonchik, and Serena Yeung. Unsupervised discovery of the long-tail in instance segmentation using hierarchical self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2603–2612, 2021. [2](#)
- [33] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020. [2](#)