Representation Tradeoffs for Hyperbolic Embeddings

Christopher De Sa[‡]   Albert Gu[†]   Christopher Ré[†]   Frederic Sala[†]

[†]Department of Computer Science, Stanford University
[‡]Department of Computer Science, Cornell University

cdesa@cs.cornell.edu, albertgu@stanford.edu, chrismre@cs.stanford.edu,fredsala@cs.stanford.edu

March 19, 2018

Abstract

Hyperbolic embeddings offer excellent quality with few dimensions when embedding hierarchical data structures like synonym or type hierarchies. Given a tree, we give a combinatorial construction that embeds the tree in hyperbolic space with arbitrarily low distortion without using optimization. On WordNet, our combinatorial embedding obtains a mean-average-precision of 0.989 with only two dimensions, while Nickel et al.'s recent construction obtains 0.87 using 200 dimensions. We provide upper and lower bounds that allow us to characterize the precision-dimensionality tradeoff inherent in any hyperbolic embedding. To embed general metric spaces, we propose a hyperbolic generalization of multidimensional scaling (h-MDS). We show how to perform exact recovery of hyperbolic points from distances, provide a perturbation analysis, and give a recovery result that allows us to reduce dimensionality. The h-MDS approach offers consistently low distortion even with few dimensions across several datasets. Finally, we extract lessons from the algorithms and theory above to design a PyTorch-based implementation that can handle incomplete information and is scalable.

1 Introduction

Recently, hyperbolic embeddings have been proposed as a way to capture hierarchy information for use in link prediction and natural language processing tasks [4, 16]. These approaches are an exciting new way to fuse rich structural information (for example, from knowledge graphs or synonym hierarchies) with the continuous representations favored by modern machine learning.

To understand the intuition behind hyperbolic embeddings' superior capacity, note that trees can be embedded with arbitrarily low distortion into the Poincaré disk, a model of hyperbolic space with only two dimensions [18]. In contrast, Bourgain's theorem [15] shows that Euclidean space is unable to obtain comparably low distortion for trees—even using an unbounded number of dimensions. Moreover, angles between embedded vectors are the same in both Euclidean and hyperbolic space (the mapping is conformal), which suggests embedded data may be easily able to integrate with downstream tasks. The optimization problem underlying hyperbolic embeddings is challenging, and as we will see it involves subtle tradeoffs.

We begin by considering the situation in which we are given an input graph that is a tree or nearly tree-like, and our goal is to produce a low-dimensional hyperbolic embedding that preserves all distances. This leads to a simple strategy that is combinatorial in that it does not minimize a surrogate loss function using gradient descent. It is both fast (nearly linear time) and has formal quality guarantees. The approach proceeds in two phases: (1) we produce an embedding of a graph into a weighted tree, and (2) we embed that tree into the hyperbolic disk. In particular, we consider an extension of an elegant embedding of trees into the Poincaré disk by Sarkar [18] and recent work on low-distortion graph embeddings into tree metrics [12]. For trees, this approach has nearly perfect quality. On the WordNet hypernym graph reconstruction, this obtains nearly perfect mean average precision (MAP) 0.989 using just two dimensions, which outperforms the best published numbers in Nickel and Kiela [16] by almost 0.12 points with 200 dimensions.

We analyze this construction to extract fundamental tradeoffs. One tradeoff involves the dimension, the properties of the graph, and the number of bits of precision - an important hidden cost. For example, on the WordNet graph, we
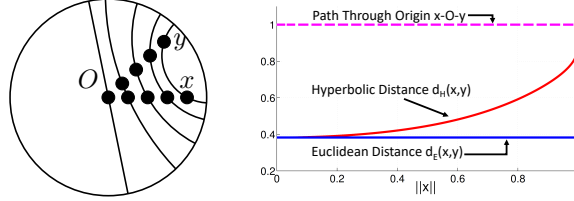
1

Figure 1: Geodesics and distances in the Poincaré disk. As x and y move towards the outside of the disk (i.e., $\|x\|, \|y\| \to 1$), the distance $d_H(x, y)$ approaches $d_H(x, O) + d_H(O, y)$.

require almost 500 bits of precision to store values from the combinatorial embedding. We can reduce this number to 32 bits, but at the cost of using 10 dimensions instead of two. We show that for a fixed precision, the dimension required scales linearly with the length of the longest path. On the other hand, the dimension scales logarithmically with the maximum degree of the tree. This suggests that hyperbolic embeddings should have high quality on hierarchies like WordNet but require large dimensions or high precision on graphs with long chains—which is supported by our experiments. A second observation is that in contrast to Euclidean embeddings, hyperbolic embeddings are not scale invariant. This motivates us to add a learnable scale term into a stochastic gradient descent-based Pytorch algorithm described below, and we show that it allows us to empirically improve the quality of embeddings.

To understand how hyperbolic embeddings perform for metrics that are far from tree-like, we consider a more general problem: given a matrix of distances that arise from points that are embeddable in hyperbolic space of dimension d (not necessarily from a graph), find a set of points that produces these distances. In Euclidean space, the problem is known as multidimensional scaling (MDS) which is solvable using PCA.[1] A key step is a transformation that effectively centers the points–without knowledge of their exact coordinates. It is not obvious how to center points in hyperbolic space, which is curved. We show that in hyperbolic space, a centering operation is possible using the Perron-Frobenius theorem. In particular, the largest eigenvalue of the distance matrix is positive and corresponds to a component-wise positive eigenvector. The components of this eigenvector allow us to define a transformation to center the points. In turn, this allows us to reduce the hyperbolic MDS problem (h-MDS) to a pair of standard eigenvalue problems, and so it can be solved with scalable power methods. Further, we extend classical perturbation analysis [19, 20]. When applied to distances from real data, h-MDS obtains low distortion on graphs that are far from tree like. However, we observe that these solution may require high precision, which is not surprising in light of our previous analysis.

Finally, we consider handling increasing amounts of noise in the model, which leads naturally into new SGD-based formulations. In traditional PCA, one may discard eigenvectors that have correspondingly small eigenvalues to cope with noise. In hyperbolic space, this approach may produce suboptimal results. Like PCA, the underlying problem is nonconvex. In contrast to PCA, the optimization problem is more challenging: the underlying problem has local minima that are not global minima. Our main technical result is that an SGD-based algorithm initialized with a h-MDS solution can recover the submanifold the data is on–even in some cases in which the data is perturbed by noise that can be full dimensional. Our algorithm essentially provides new recovery results for convergence for Principal Geodesic Analysis (PGA) in hyperbolic space [9].

All of our results can handle incomplete distance information through standard techniques. Using the observations above, we implemented an SGD algorithm that minimizes the loss derived from the PGA loss using PyTorch.[2]

2 Background

We provide intuition connecting hyperbolic space and tree distances, discuss the metrics used to measure embedding fidelity, and provide the relationship between reconstruction and learning for graph embeddings.

---

[1] There is no perfect analog of PCA in hyperbolic space [17].

[2] A minor instability with Chamberlain et al. [4], Nickel and Kiela [16]'s formulation is that one must guard against NANs. This instability may be unavoidable in formulations that minimize hyperbolic distance with gradient descent, as the derivative of the hyperbolic distance has a singularity, that is, $\lim_{y \to x} \partial_x |d_H(x, y)| \to \infty$ for any $x \in \mathbb{H}$ in which d_H is the hyperbolic distance function. This issue can be mitigated by minimizing d_H^2, which does have a continuous derivative throughout $\mathbb{H}$. We propose to do so in Section 4.2 and discuss this further in the Appendix.

Hyperbolic spaces The Poincaré disk $\mathbb{H}_2$ is a two-dimensional model of hyperbolic geometry with points located in the interior of the unit disk, as shown in Figure 1. A natural generalization of $\mathbb{H}_2$ is the Poincaré ball $\mathbb{H}_r$, with elements inside the unit ball. The Poincaré models offer several useful properties, chief among which is mapping conformally to Euclidean space. That is, angles are preserved between hyperbolic and Euclidean space. Distances, on the other hand, are not preserved, but are given by

$$d_H(x,y) = \operatorname{acosh}\left(1 + 2\frac{\|x-y\|^2}{(1-\|x\|^2)(1-\|y\|^2)}\right).$$

There are some potentially unexpected consequences of this formula, and a simple example gives intuition about a key technical property that allows hyperbolic space to embed trees. Consider three points: the origin 0, and points x and y with $\|x\| = \|y\| = t$ for some $t > 0$. As shown on the right of Figure 1, as $t \to 1$ (i.e., the points move towards the outside of the disk), in flat Euclidean space, the ratio $\frac{d_E(x,y)}{d_E(x,0)+d_E(0,y)}$ is *constant* with respect to t (blue curve). In contrast, the distance $d_H(x,y)$ approaches $d_H(x,0) + d_H(0,y)$ (red and pink curves). That is, the shortest path between x and y is almost the same as the path through the origin. This is analogous to the property of trees in which the shortest path between two sibling nodes is the path through their parent. This tree-like nature of hyperbolic space is the key property exploited by embeddings. Moreover, this property holds for arbitrarily small angles between x and y.

Lines and geodesics There are two types of geodesics (shortest paths) in the Poincaré disk model of hyperbolic space: segments of circles that are orthogonal to the disk surface, and disk diameters [2]. Our algorithms and proofs make use of a simple geometric fact: *isometric* reflection across geodesics (preserving hyperbolic distances) is represented in this Euclidean model as a *circle inversion*. A particularly important reflection associated with each point x is the one that takes x to the origin [2, p. 268].

Embeddings and fidelity measures An *embedding* is a mapping $f : U \to V$ for spaces U, V with distances d_U, d_V. We measure the quality of embeddings with several *fidelity measures*, presented here from most local to most global.

Recent work [16] proposes using the *mean average precision* (MAP). For a graph $G = (V, E)$, let $a \in V$ have neighborhood $\mathcal{N}_a = \{b_1, b_2, \ldots, b_{\deg(a)}\}$, where $\deg(a)$ denotes the degree of a. In the embedding f, consider the points closest to $f(a)$, and define R_{a,b_i} to be the smallest set of such points that contains b_i (that is, R_{a,b_i} is the smallest set of nearest points required to retrieve the ith neighbor of a in f). Then, the MAP is defined to be

$$\text{MAP}(f) = \frac{1}{|V|}\sum_{a \in V}\frac{1}{|\mathcal{N}_a|}\sum_{i=1}^{|\mathcal{N}_a|}\text{Precision}(R_{a,b_i}) = \frac{1}{|V|}\sum_{a \in V}\frac{1}{\deg(a)}\sum_{i=1}^{|\mathcal{N}_a|}\frac{|\mathcal{N}_a \cap R_{a,b_i}|}{|R_{a,b_i}|}.$$

We have $\text{MAP}(f) \leq 1$, with equality as the best case. Note that MAP is not concerned with the underlying distances at all, but only the ranks between the distances of immediate neighbors. It is a *local* metric.

The standard metric for graph embeddings is distortion D. For an n point embedding,

$$D(f) = \frac{1}{\binom{n}{2}}\left(\sum_{u,v \in U : u \neq v}\frac{|d_V(f(u),f(v)) - d_U(u,v)|}{d_U(u,v)}\right).$$

The best distortion is $D(f) = 0$, preserving the edge lengths exactly. This is a *global* metric, as it depends directly on the underlying distances rather than the local relationships between distances. A variant of this, the worst-case distortion D_{wc}, is the metric defined by

$$D_{\text{wc}}(f) = \frac{\max_{u,v \in U : u \neq v} d_V(f(u),f(v))/d_U(u,v)}{\min_{u,v \in U : u \neq v} d_V(f(u),f(v))/d_U(u,v)}.$$

That is, the wost-case distortion is the ratio of the maximal expansion and the minimal contraction of distances. Note that scaling the unit distance does not affect D_{wc}. The best worst-case distortion is $D_{\text{wc}}(f) = 1$.

The intended application of the embedding informs the choice of metric. For applications where the underlying distances are important, distortion is useful. On the other hand, if only rankings matter, MAP may suffice. This choice is important: as we shall see, different embedding algorithms implicitly target different metrics.

Algorithm 1 Sarkar's Construction

1: **Input:** Node a with parent b, children to place $c_1, c_2, \ldots, c_{\deg(a)-1}$, partial embedding f containing an embedding for a and b, scaling factor τ
2: $(0, z) \leftarrow \text{reflect}_{f(a) \to 0}(f(a), f(b))$
3: $\theta \leftarrow \arg(z)$ {angle of z from x-axis in the plane}
4: **for** $i \in \{1, \ldots, \deg(a) - 1\}$ **do**
5:    $y_i \leftarrow \left(\frac{e^{\tau}-1}{e^{\tau}+1} \cdot \cos\left(\theta + \frac{2\pi i}{\deg(a)}\right), \frac{e^{\tau}-1}{e^{\tau}+1} \cdot \sin\left(\theta + \frac{2\pi i}{\deg(a)}\right)\right)$
6: **end for**
7: $(f(a), f(b), f(c_1), \ldots, f(c_{\deg(a)-1})) \leftarrow \text{reflect}_{0 \to f(a)}(0, z, y_1, \ldots, y_{\deg(x)-1})$
8: **Output:** Embedded $\mathbb{H}_2$ vectors $f(c_1), f(c_2), \ldots, f(c_{\deg(a)-1})$

Reconstruction and learning In the case where we lack a full set of distances, we can deal with the missing data in one of two ways. First, we can use the triangle inequality to recover the missing distances. Second, we can access the scaled Euclidean distances (the inside of the acosh in $d_H(x, y)$), and then the resulting matrix can be recovered with standard matrix completion techniques [3]. Afterwards, we can proceed to compute an embedding using any of the approaches discussed in this paper. We quantify the error introduced by this process experimentally in Section 5.

3 Combinatorial Constructions

We first focus on hyperbolic tree embeddings—a natural approach considering the tree-like behavior of hyperbolic space. We review the embedding of Sarkar [18] to higher dimensions. We then provide novel analysis about the precision of the embeddings that reveals fundamental limits of hyperbolic embeddings. In particular, we characterize the bits of precision needed for hyperbolic representations. We then extend the construction to r dimensions, and we propose to use Steiner nodes to better embed general graphs as trees building on a condition from I. Abraham et al. [12].

Embedding trees The nature of hyperbolic space lends itself towards excellent tree embeddings. In fact, it is possible to embed trees into the Poincaré disk $\mathbb{H}_2$ with arbitrarily low distortion [18]. Remarkably, trees cannot be embedded into Euclidean space with arbitrarily low distortion for *any* number of dimensions. These notions motivate the following two-step process for embedding hierarchies into hyperbolic space.

1. Embed the graph $G = (V, E)$ into a tree T,

2. Embed T into the Poincaré ball $\mathbb{H}_d$.

We refer to this process as the *combinatorial construction*. Note that we are not required to minimize a loss function. We begin by describing the second stage, where we extend an elegant construction from Sarkar [18].

3.1 Sarkar's Construction

Algorithm 1 implements a simple embedding of trees into $\mathbb{H}_2$. The algorithm takes as input a scaling factor τ a node a (of degree $\deg(a)$) from the tree with parent node b. Suppose a and b have already been embedded into $\mathbb{H}_2$ and have corresponding embedded vectors $f(a)$ and $f(b)$. The algorithm places the children $c_1, c_2, \ldots, c_{\deg(a)-1}$ into $\mathbb{H}_2$ through a two-step process.

First, $f(a)$ and $f(b)$ are reflected across a geodesic (using circle inversion) so that $f(a)$ is mapped onto the origin 0 and $f(b)$ is mapped onto some point z. Next, we place the children nodes to vectors $y_1, \ldots, y_{d-1}$ equally spaced around a circle with radius $\frac{e^{\tau}-1}{e^{\tau}+1}$ (which is a circle of radius τ in the hyperbolic metric), and maximally separated from the reflected parent node embedding z. Lastly, we reflect all of the points back across the geodesic. Note that the isometric properties of reflections imply that all children are now at hyperbolic distance exactly τ from $f(a)$.

To embed the entire tree, we place the root at the origin O and its children in a circle around it (as in Step 5 of Algorithm 1), then recursively place their children until all nodes have been placed. Notice this construction runs in linear time.

3.2 Analyzing Sarkar's Construction

The *Voronoi cell* around a node $a \in T$ consists of points $x \in \mathbb{H}_2$ such that $d_H(f(a), x) \leq d_H(f(b), x)$ for all $b \in T$ distinct from a. That is, the cell around a includes all points closer to $f(a)$ than to any other embedded node of the tree. Sarkar's construction produces Delauney embeddings: embeddings where the Voronoi cells for points a and b touch only if a and b are neighbors in T. Thus this embedding will preserve neighborhoods.

A key technical idea exploited by Sarkar [18] is to scale all the edges by a factor τ before embedding. We can then recover the original distances by dividing by τ. This transformation exploits the fact that hyperbolic space is not *scale invariant*. Sarkar's construction always captures neighbors perfectly, but Figure 1 implies that increasing the scale preserves the distances between farther nodes better. Indeed, if one sets $\tau = \frac{1+\varepsilon}{\varepsilon}\left(2\log\frac{\deg_{\max}}{\pi/2}\right)$, then the worst-case distortion D of the resulting embedding is no more than $1 + \varepsilon$. For trees, Sarkar's construction has arbitrarily high fidelity. However, this comes at a cost: the scaling τ affects the bits of precision required. In fact, we will show that the precision scales logarithmically with the degree of the tree—but linearly with the maximum path length. We use this to better understand the situations in which hyperbolic embeddings obtain high quality.

How many bits of precision do we need to represent points in $\mathbb{H}_2$? If $x \in \mathbb{H}_2$, then $\|x\| < 1$, so we need sufficiently many bits so that $1 - \|x\|$ will not be rounded to zero. This requires roughly $-\log(1 - \|x\|) = \log\frac{1}{1-\|x\|}$ bits. Say we are embedding two points x, y at distance d. As described in the background, there is an isometric reflection that takes a pair of points (x, y) in $\mathbb{H}_2$ to $(0, z)$ while preserving their distance, so without loss of generality we have that

$$d = d_H(x, y) = d_H(0, z) = \operatorname{acosh}\left(1 + 2\frac{\|z\|^2}{1 - \|z\|^2}\right).$$

Rearranging the terms, we have

$$\frac{\cosh(d) + 1}{2} = \frac{1}{1 - \|z\|^2} \geq \frac{1/2}{1 - \|z\|}.$$

Thus, the number of bits we want so that $1 - \|z\|$ will not be rounded to zero is $\log(\cosh(d) + 1)$. Since $\cosh(d) = (\exp(d) + \exp(-d))/2$, this is roughly d bits. That is, in hyperbolic space, we need about d bits to express distances of d (rather than $\log d$ as we would in Euclidean space). This result will be of use below.

Now we consider the largest distance in the embeddings produced by Algorithm 1. If the longest path in the tree is ℓ, and each edge has length $\tau = \frac{1+\varepsilon}{\varepsilon}\left(2\log\frac{\deg_{\max}}{\pi/2}\right)$, the largest distance is $O(\frac{1+\varepsilon}{\varepsilon}\ell\log\deg_{\max})$, and we require this number of bits for the representation.

We interpret this expression. Note that $\deg_{\max}$ is inside the $\log$ term, so that a bushy tree is not penalized much in precision. On the other hand, the longest path length ℓ is not, so that hyperbolic embeddings struggle with long paths. Moreover, by selecting an explicit graph, we derive a matching lower bound, concluding that to achieve a distortion ε, any construction requires $\Omega\left(\frac{1}{\varepsilon}\ell\right)$ bits, which nearly matches the upper bound of the combinatorial construction. The argument follows from selecting a graph consisting of $4n + 1$ nodes in a tree with a single root and four equal length chains. We show that preserving the distances between the nodes at the end of each chain requires $\Omega(n\varepsilon^{-1})$ bits. The proof of this result is described in Appendix D.

3.3 Improving the Construction

Our next contribution is a generalization of the construction from the disk $\mathbb{H}_2$ to the ball $\mathbb{H}_r$. Our construction follows the same line as Algorithm 1, but since we have r dimensions, the step where we place children spaced out on a circle around their parent now uses a hypersphere.
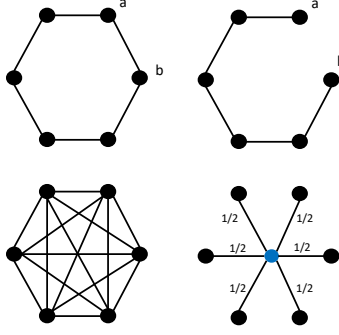
Figure 2: Top. Cycles are a challenge for tree embeddings: $d_G(a, b)$ goes from 1 to 5. Bottom. Steiner nodes can help: adding a node (blue) and weighting edges maintains the pairwise distances.

Spacing out points on the hypersphere is a classic problem known as *spherical coding* [6]. As we shall see, the number of children that we can place for a particular angle grows with the dimension. Since the required scaling factor τ gets larger for as the angle decreases, we can reduce τ for a particular embedding by increasing the dimension. Note that increasing the dimension helps with bushy trees (large $\deg_{\max}$), but has limited effect on tall trees with small $\deg_{\max}$. We show

Proposition 3.1. *The generalized $\mathbb{H}_r$ combinatorial construction has distortion at most $1 + \varepsilon$ and requires at most $O(\frac{1+\varepsilon}{\varepsilon} \frac{\ell}{r} \log \deg_{\max})$ bits to represent a node component for $r \leq (\log \deg_{\max}) + 1$, and $O(\frac{1+\varepsilon}{\varepsilon} \ell)$ bits for $r > (\log \deg_{\max}) + 1$.*

The algorithm for the generalized $\mathbb{H}_r$ combinatorial construction replaces Step 5 in Algorithm 1 with a node placement step based on ideas from coding theory. The children are placed at the vertices of a hypercube inscribed into the unit hypersphere (and afterwards scaled by τ). Each component of a hypercube vertex has the form $\frac{\pm 1}{\sqrt{r}}$. We index these points using binary sequences $a \in \{0, 1\}^r$ in the following way:

$$x_a = \left(\frac{(-1)^{a_1}}{\sqrt{r}}, \frac{(-1)^{a_2}}{\sqrt{r}}, \ldots, \frac{(-1)^{a_r}}{\sqrt{r}} \right).$$

We can space out the children by controlling the distances between the children. This is done in turn by selecting a set of binary sequences with a prescribed minimum Hamming distance—a binary error-correcting code—and placing the children at the resulting hypercube vertices. We provide more details on this technique and our choice of code in the appendix.

3.4 Embedding into Trees

We revisit the first step of the construction: embedding graphs into trees. There are fundamental limits to how well graphs can be embedded into trees; in general, breaking long cycles inevitably adds distortion, as shown in Figure 2. We are inspired by a measure of this limit, the δ-*4 points condition* introduced in I. Abraham et al. [12]. A graph on n nodes that satisfies the δ-4 points condition has distortion at most $(1 + \delta)^{c_1 \log n}$ for some constant c_1. This result enables our end-to-end embedding to achieve a distortion of at most

$$D(f) \leq (1 + \delta)^{c_1 \log n}(1 + \varepsilon).$$

The result in I. Abraham et al. [12] builds a tree with *Steiner* nodes. These additional nodes can help control the distances in the resulting tree.

Example 3.1. *Embed a complete graph on $\{1, 2, \ldots, n\}$ into a tree. The tree will have a central node, say 1, w.l.o.g., connected to every other node; the shortest paths between pairs of nodes in $\{2, \ldots, n\}$ go from distance 1 in the graph to distance 2 in the tree. However, we can introduce a Steiner node $n + 1$ and connect it to all of the nodes, with edge weights of $\frac{1}{2}$. This is shown in Figure 2. The distance between any pair of nodes in $\{1, \ldots, n\}$ remains 1.*

6

Note that introducing Steiner nodes can produce a weighted tree, but Algorithm 1 readily extends to the case of weighted trees by modifying Step 5. We propose using the Steiner tree algorithm in I. Abraham et al. [12] (used to achieve the distortion bound) for real embeddings, and we rely on it for our experiments in Section 5. In summary, the key takeaways of our analysis in this section are:

- There is a fundamental tension between precision and quality in hyperbolic embeddings.

- Hyperbolic embeddings have an exponential advantage in space compared to Euclidean embeddings for short, bushy hierarchies, but will have less of an advantage for graphs that contain long paths.

- Choosing an appropriate scaling factor τ is critical for quality. Later, we will propose to learn this scale factor automatically for computing embeddings in PyTorch.

- Steiner nodes can help improve embeddings of graphs.

4   Hyperbolic Multidimensional Scaling

In this section, we explore a fundamental and more general question than the previous section: if we are given the pairwise distances arising from a set of points in hyperbolic space, can we recover the points? The equivalent problem for Euclidean distances is solved with multidimensional scaling (MDS). The goal of this section is to propose *hyperbolic MDS* (h-MDS). We describe and overcome the additional technical challenges imposed by hyperbolic distances and show that, remarkably, exact recovery is possible if the dimension of the points is known. Afterwards we propose a technique for dimensionality reduction using principal geodesics analysis (PGA).

4.1   Exact Hyperbolic MDS

Suppose that there is a set of hyperbolic points $x_1, \ldots, x_n \in \mathbb{H}_r$, embedded in the Poincaré sphere and written $X \in \mathbb{R}^{n \times r}$ in matrix form. We are given all the pairwise distances $d_{i,j} = d_H(x_i, x_j)$. We observe $d_{i,j}$ but do not observe X: our goal is use the observed $d_{i,j}$'s to recover X (or some other set of points with the same pairwise distances $d_{i,j}$). First, we rewrite the observations as

$$Y_{i,j} = \frac{\cosh(d_{i,j}) - 1}{2} = \frac{\|x_i - x_j\|^2}{(1 - \|x_i\|^2)(1 - \|x_j\|^2)}. \tag{1}$$

Next, let $S \in \mathbb{R}^{n \times n}$ be the diagonal matrix where $S_{i,i} = 1/(1 - \|x_i\|^2)$ and let $v \in \mathbb{R}^n$ be the vector where by $v_i = \|x_i\|^2$. Then, we can write (1) in matrix form as

$$Y = S(1v^T + v1^T)S - 2SXX^TS \tag{2}$$

This form reminds us of the Euclidean version of the problem, where the observations are simply $\|x_i - x_j\|^2$: there, the Euclidean MDS algorithm proceeds by projecting out the $1v^T$ terms and using PCA to recover X from XX^T. Here, however, we face the additional difficulty of the $S_{i,i} = 1/(1 - \|x_i\|^2)$ terms, equivalent to the norms of the points, which we *do not* have access to. Nevertheless, we show that exact recovery is still possible.

Algorithm 2 is our complete algorithm, and for the remainder of this section we will describe how and why it works. We first provide a technical result that simplifies the algorithm by effectively centering the points in hyperbolic space. Next, we explain the key relationships that enable us to find S and v. Finally, we reduce the problem to an Euclidean PCA problem, and we show how we can compute the embeddings.

Easy Recovery   Notice that Y has nonnegative entries, and so Perron's theorem allows us to conclude that the principal eigenvalue of Y is positive, and that the principal eigenvector is componentwise non-negative [11]. We call such a principal eigenvector $\hat{u}$. We show in the Appendix that if a distance matrix $d_{i,j}$ is derived from a set of hyperbolic points, it has an exact embedding X that satisfies the *easy recovery property*: $X^T S \hat{u} = 0$. For this exact embedding, $2SXX^T S \hat{u} = 0$, and so $\hat{u}$ must also be an eigenvector of $Z = S(1v^T + v1^T)S$ alone. As we shall see, we can use this

---

Algorithm 2 h-MDS

1: **Input: Distance matrix $d_{i,j}$ and rank** r
2: Compute scaled distances $Y_{i,j} = \frac{\cosh(d_{i,j}) - 1}{2}$
3: Get principal unit eigenvector $\hat{u}$ and eigenvalue λ of Y
4: $\beta \leftarrow 1 + \frac{(1^T \hat{u})^2}{\lambda \hat{u}^T \hat{u}}$
5: $\alpha \leftarrow \beta - \sqrt{\beta^2 - 1}$
6: Scale u: $u \leftarrow \frac{\hat{u}}{1^T \hat{u}} \lambda (1 - \alpha)$
7: Retrieve scaling matrix $S \leftarrow \text{diag}\left(\frac{u+1}{1+\alpha}\right)$
8: Compute norms $v \leftarrow S^{-1}\left(\frac{u - \alpha 1}{1 + \alpha}\right)$
9: $XX^T \leftarrow \frac{1}{2}(v1^T + 1v^T - S^{-1}YS^{-1})$
10: **return** $\text{PCA}(XX^T, r)$

---

fact to recover S given $\hat{u}$. We can compute $\hat{u}$ given Y using any principal eigensolver (in our implementation a standard power method).

Computing Eigenvectors   Because Z is a rank-2 matrix of the form $xy^T + yx^T$, it is easy to see that both eigenvectors of Z are of the form $u = Sv + \alpha S1$ for some $\alpha \in \mathbb{R}$. If we knew α together with u, it would be possible to recover v and S to complete the problem. The challenge is that we know the direction, which is $\hat{u}$, but not the length of the eigenvector $u = Sv + \alpha S1$. Our idea is to find an equation that involves α and depends on u in a *scale invariant* way—so that we can use $\hat{u}$, which we have access to, instead of u, whose scale is unknown to us. Formally, we observe that the following quadratic equation holds, which we derive in the Appendix.

$$\frac{(1 - \alpha)^2}{\alpha} = \frac{2}{\lambda} \frac{(1^T u)^2}{u^T u}. \tag{3}$$

This equation satisfies our goal: the quantity on the right-hand side is invariant to the scale of u, and so we can compute it from $\hat{u}$. But the quantity is quadratic in α which means there may be more than one solution. We define $\beta := 1 + \frac{(1^T u)^2}{\lambda u^T u}$ and pick the positive root $\alpha = \beta - \sqrt{\beta^2 - 1}$ of the resulting quadratic. An inspection shows that this is the correct root, and we prove this in Appendix E.

Centering XX^T   Given u and α, we can compute Sv, $S1$, and then v. We can then compute XX^T in terms of Y by starting from (2), multiplying both sides by S^{-1}, and subtracting off the $1v^T + v1^T$ terms. This is analogous to the centering operation from Euclidean MDS, which centers using the operator $I - \frac{1}{n}11^T$. As is standard, the principal components of XX^T now are our embedding vectors into the hyperbolic space.

Algorithm 2 gives the full description of how to compute these embeddings exactly using standard PCA solvers. In the appendix, we derive a perturbation analysis of this algorithm following Sibson [19], which exactly captures the worst-case perturbation sensitivity. As with PCA, the perturbation sensitivity scales linearly in the smallest eigenvalue and the number of entries—but it is exponentially more sensitive to perturbations of elements with large norms, as one might expect.

4.2   Reducing Dimensionality with PGA

Sometimes we are given a high-rank embedding (resulting from h-MDS, for example), and wish to find a lower-rank version. In Euclidean space, one can get the optimal lower rank embedding by simply discarding components. However, this may not be the case in hyperbolic space. Motivated by this, we study dimensionality reduction in hyperbolic space.

As hyperbolic space does not have a linear subspace structure like Euclidean space, we need to define what we mean by lower-dimensional. We follow Principal Geodesic Analysis [9]. Consider an initial embedding with points $x_1, \ldots, x_n \in \mathbb{H}_2$ and let $d_H : \mathbb{H}_2 \times \mathbb{H}_2 \to \mathbb{R}_+$ be the hyperbolic distance. Suppose we want to map this embedding
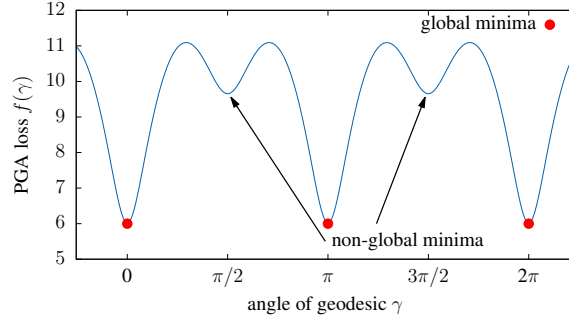
Figure 3: The PGA objective of an example task where the input dataset in the Poincaré disk is $x_1 = (0.8, 0)$, $x_2 = (-0.8, 0)$, $x_3 = (0, 0.7)$ and $x_4 = (0, -0.7)$. Note the presence of non-optimal local minima, unlike PCA.

onto a one-dimensional subspace. (Note that we are considering a two-dimensional embedding and one-dimensional subspace here for simplicity, and these results immediately extend to higher dimensions.) In this case, the goal of PGA is to find a geodesic $\gamma : [0, 1] \to \mathbb{H}_2$ that passes through the mean of the points and that minimizes the squared error (or variance):

$$f(\gamma) = \sum_{i=1}^{n} \min_{t \in [0,1]} d_H(\gamma(t), x_i)^2.$$

This expression can be simplified significantly and reduced to a minimization in Euclidean space. First, we find the mean of the points, the point $\bar{x}$ which minimizes $\sum_{i=1}^{n} d_H(\bar{x}, x_i)^2$; this definition in terms of distances generalizes the mean in Euclidean space.[3] Next, we reflect all the points x_i so that their mean is 0 in the Poincaré disk model; we can do this using a circle inversion that maps $\bar{x}$ onto 0. In the Poincaré disk model, a geodesic through the origin is a Euclidean line, and the action of the reflection across this line is the same in both Euclidean and hyperbolic space. Coupled with the fact that reflections are isometric, if γ is a line through 0 and R_γ is the reflection across γ, we have

$$d_H(\gamma, x) = \min_{t \in [0,1]} d_H(\gamma(t), x) = \frac{1}{2} d_H(R_l x, x).$$

Combining this with the Euclidean reflection formula and the hyperbolic metric produces

$$f(\gamma) = \frac{1}{4} \sum_{i=1}^{n} \operatorname{acosh}^2 \left(1 + \frac{8 d_E(\gamma, x_i)^2}{(1 - \|x_i\|^2)^2} \right),$$

in which d_E is the Euclidean distance from a point to a line. If we define $w_i = \sqrt{8} x_i / (1 - \|x_i\|^2)$ this reduces to the simplified expression

$$f(\gamma) = \frac{1}{4} \sum_{i=1}^{n} \operatorname{acosh}^2 \left(1 + d_E(\gamma, w_i)^2 \right).$$

Notice that *the loss function is not convex*. We observe that there can be multiple local minima that are attractive and stable, in contrast to PCA. Figure 3 illustrates this nonconvexity on a simple dataset in $\mathbb{H}_2$ with only four examples. This makes globally optimizing the objective difficult.

Nevertheless, there will always be a region Ω containing a global optimum γ^* that is convex and admits an efficient projection, and where f is convex when restricted to Ω. Thus it is possible to build a gradient descent-based algorithm to recover lower-dimensional subspaces: for example, we built a simple optimizer in PyTorch. We also give a sufficient condition on the data for f above to be convex.

Lemma 4.1. *For hyperbolic PGA if for all i,*

$$\operatorname{acosh}^2 \left(1 + d_E(\gamma, w_i)^2 \right) < \min \left(1, \frac{1}{3} \|w_i\|^2 \right)$$

then f is locally convex at γ.

---

[3] As we noted earlier, considering the distances without squares leads to a non-continuously-differentiable formulation.

9

Dataset	Nodes	Edges	Comment
Bal. Tree	40	39	Tree
Phy. Tree	344	343	Tree
CS PhDs	1025	1043	Tree-like
WordNet*	74374	75834	Tree-like
Diseases	516	1188	Dense

Table 1: Datasets Statistics.

As a result, if we initialize in and optimize over a region that contains γ^* and where the condition of Lemma 4.1 holds, then gradient descent will be guaranteed to converge to γ^*. We can turn this result around and read it as a recovery result: if the noise is bounded in this regime, then we are able to provably recover the correct low-dimensional embedding.

5 Experiments

Dataset	C-$\mathbb{H}_2$	FB $\mathbb{H}_2$	h-MDS	PyTorch	PWS	PCA	FB
Bal. Tree	**0.013**	0.425	**0.077**	0.034	0.020	0.496	0.236
Phy. Tree	**0.006**	0.832	**0.039**	0.237	0.092	0.746	0.583
CS PhDs	**0.286**	0.542	**0.149**	0.298	0.187	0.708	0.336
Diseases	**0.147**	0.410	0.111	**0.080**	0.108	0.595	0.764

Table 2: Distortion measures using combinatorial and h-MDS techniques, compared against PCA and results from Nickel and Kiela [16]. Closer to 0 is better.

Dataset	C-$\mathbb{H}_2$	FB $\mathbb{H}_2$	h-MDS	PyTorch	PWS	PCA	FB
Bal. Tree	**1.0**	0.846	**1.0**	**1.0**	**1.0**	**1.0**	0.859
Phy. Tree	**1.0**	0.718	0.675	0.951	0.998	**1.0**	0.811
CS PhDs	**0.991**	0.567	0.463	0.799	**0.945**	0.541	0.78
Diseases	**0.822**	0.788	0.949	0.995	0.897	**0.999**	0.934

Table 3: MAP measures using combinatorial and h-MDS techniques, compared against PCA. Closer to 1 is better.

We evaluate the proposed approaches and compare against existing methods. We hypothesize that for tree-like data, the combinatorial construction offers the best performance. For general data, we expect h-MDS to produce the lowest distortion, while it may have low MAP due to precision limitations. We anticipate that dimension is a critical factor (outside of the combinatorial construction). Finally, we expect that the MAP of h-MDS techniques can be improved by learning the correct scale and weighting the loss function as suggested in earlier sections. In the Appendix, we report on additional datasets, parameters found by the combinatorial construction, the effect of hyperparameters, and on reducing the precision required by the combinatorial construction by increasing the dimension; as we observed in Section 3, we can reduce the scaling by a factor depending on $\deg_{\max}$, but the $\frac{1+\epsilon}{\epsilon}$ term is significant. For $\varepsilon = 1.0$, however, we can store components in 64 bits.

Datasets We consider trees, tree-like hierarchies, and graphs that are not tree-like. First, we consider hierarchies that form trees: fully-balanced trees along with phylogenetic trees expressing genetic heritage. Similarly, we used hierarchies that are nearly tree-like: WordNet hypernym (the largest connected component from Nickel and Kiela [16]) and a graph of Ph.D. advisor-advisee relationships. Also included are datasets that vary in their tree nearness, such as biological sets involving disease relationships and protein interactions in yeast bacteria.
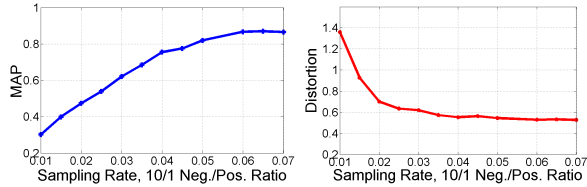
Figure 4: Learning from incomplete information. The distance matrix is sampled, completed, and embedded.

Rank	No Scale	Learned Scale	Exp. Weighting
50	0.481	0.508	**0.775**
100	0.688	0.681	**0.882**
200	0.894	0.907	**0.963**

Table 4: Ph.D. dataset. Improved MAP performance of PyTorch implementation using a modified PGA-like loss function.

Approaches   Combinatorial embeddings into $\mathbb{H}_2$ are done using Steiner trees generated from a randomly selected root for the $\varepsilon = 0.1$ precision setting; others are considered in the Appendix. We performed h-MDS in floating point precision. We also include results for our PyTorch implementation of an SGD-based algorithm (described later), as well as a warm start version initialized with the high-dimensional combinatorial construction. We compare against classical MDS (i.e., PCA), and the optimization-based approach Nickel and Kiela [16], which we call FB. The experiments for h-MDS, PyTorch SGD, PCA, and FB used dimensions of 2,5,10,50,100,200; we recorded the best resulting MAP and distortion. Due to the large scale, we did not replicate the best FB numbers on large graphs (e.g., WordNet). As a result, we report their best published numbers for comparison. For the WordNet graph we use a random BFS tree rather than a Steiner tree.

Quality   In Table 2, we report the distortion. As expected, when the graph is a tree or tree-like the combinatorial construction has exceedingly low distortion. Because h-MDS is meant to recover points exactly, we hypothesized that h-MDS would offer very low distortion on these datasets. We confirm this hypothesis: among h-MDS, PCA, and FB, h-MDS consistently offers the best (lowest) distortion, producing, for example, a distortion of 0.039 on the phylogenetic tree dataset.

Table 3 reports the MAP measure, which is a local measure. We expect that the combinatorial construction performs well for tree-like hierarchies. This is indeed the case: on trees and tree-like graphs, the MAP is close to 1, improving on approaches such as FB that rely on optimization. On larger graphs like WordNet, our approach yields a MAP of 0.989–improving on the FB MAP result of 0.870 at 200 dimensions. This is exciting because the combinatorial approach is deterministic and linear-time. In addition, it suggests that this refined understanding of hyperbolic embeddings may be used to improve the quality and runtime state of the art constructions. As expected, the MAP of the combinatorial construction decreases as the graphs are less tree-like. Interestingly, h-MDS solved in floating point indeed struggles with MAP. We separately confirmed that it is indeed due to precision using a high-precision solver, which obtains a perfect MAP—but uses 512 bits of precision. It may be possible to compensate for this with scaling, but we did not explore this possibility.

SGD-Based Algorithm   We also built an SGD-based algorithm implemented in PyTorch. Here the loss function is equivalent to the PGA loss, and so is continuously differentiable. We use this to verify two claims:

Learned Scale. In Table 4, we verify the importance of scaling that our analysis suggests; our implementation has a simple learned scale parameter. Moreover, we added an exponential weighting to the distances in order to penalize long paths, thus improving the local reconstruction. These techniques indeed improve the MAP; in particular, the learned scale provides a better MAP at lower rank. We hope these techniques can be useful in other embedding techniques.

Incomplete Information. To evaluate our algorithm's ability to deal with incomplete information, we examine the

quality of recovered solutions as we sample the distance matrix. We set the sampling rate of non-edges to edges at $10:1$ following Nickel and Kiela [16]. We examine the phylogenetic tree, which is full rank in Euclidean space. In Figure 4, we are able to recover a good solution with a small fraction of the entries for the phylogenetic tree dataset; for example, we sampled approximately 4% of the graph but provide a MAP of 0.74 and distortion of less than 0.6. Understanding the sample complexity for this problem is an interesting theoretical question.

6 Conclusion and Future Work

Hyperbolic embeddings embed hierarchical information with high fidelity and few dimensions. We explored the limits of this approach by describing scalable, high quality algorithms. We hope the techniques here encourage more follow-on work on the exciting techniques of Chamberlain et al. [4], Nickel and Kiela [16]. As future work, we hope to explore how hyperbolic embeddings can be most effectively incorporated into downstream tasks such as LSTMs.

References

[1] M. Abu-Ata and F. F. Dragan. Metric tree-like structures in real-world networks: an empirical study. *Networks*, 67 (1):49–68, 2015.

[2] D. Brannan, M. Esplen, and J. Gray. *Geometry*. Cambridge University Press, Cambridge, UK, 2012.

[3] E. Candes and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

[4] B. P. Chamberlain, J. R. Clough, and M. P. Deisenroth. Neural embeddings of graphs in hyperbolic space. *arXiv preprint, arXiv:1705.10359*, 2017.

[5] W. Chen, W. Fang, G. Hu, and M. W. Mahoney. On the hyperbolicity of small-world and tree-like random graphs. In *International Symposium on Algorithms and Computation (ISAAC) 2012*, pages 278–288, Taipei, Taiwan, 2012.

[6] J. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*. Springer, New York, NY, 1999.

[7] A. Cvetkovski and M. Crovella. Hyperbolic embedding and routing for dynamic graphs. In *IEEE INFOCOM 2009*, 2009.

[8] D. Eppstein and M. Goodrich. Succinct greedy graph drawing in the hyperbolic plane. In *Proc. of the International Symposium on Graph Drawing (GD 2011)*, pages 355–366, Eindhoven, Netherlands, 2011.

[9] P. Fletcher, C. Lu, S. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005, 2004.

[10] M. Gromov. Hyperbolic groups. In *Essays in group theory*. Springer, 1987.

[11] R. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 2013.

[12] I. Abraham et al. Reconstructing approximate tree metrics. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of Distributed Computing (PODC)*, pages 43–52, Portland, Oregon, 2007.

[13] M. Jenssen, F. Joos, and W. Perkin. On kissing numbers and spherical codes in high dimensions. *arXiv preprint, arXiv:1803.02702*, 2018.

[14] R. Kleinberg. Geographic routing using hyperbolic space. In *26th IEEE International Conference on Computer Communications (ICC)*, pages 1902–1909, 2007.

[15] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[16] M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, Long Beach, CA, 2017.

[17] X. Pennec. Barycentric subspace analysis on manifolds. *Annals of Statistics*, to appear 2017.

[18] R. Sarkar. Low distortion Delaunay embedding of trees in hyperbolic plane. In *Proc. of the International Symposium on Graph Drawing (GD 2011)*, pages 355–366, Eindhoven, Netherlands, 2011.

[19] R. Sibson. Studies in the robustness of multidimensional scaling: Procrustes statistics. *Journal of the Royal Statistical Society, Series B*, 40(2):234–238, 1978.

[20] R. Sibson. Studies in the robustness of multidimensional scaling: Perturbational analysis of classical scaling. *Journal of the Royal Statistical Society, Series B*, 41(2):217–229, 1979.

[21] M. Zhang and P. Fletcher. Probabilistic principal geodesic analysis. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, Lake Tahoe, NV, 2013.

A   Glossary of Symbols

Symbol	Used for
x, y, z	vectors in the Poincaré ball model of hyperbolic space
d_H	metric distance between two points in hyperbolic space
d_E	metric distance between two points in Euclidean space
d_U	metric distance between two points in metric space U
d	a particular distance value
$d_{i,j}$	the distance between the ith and jth points in an embedding
$\mathbb{H}_r$	the Poincaré ball model of r-dimensional Hyperbolic space
r	the dimension of a Hyperbolic space
$\mathbb{H}$	Hyperbolic space of an unspecified or arbitrary dimension
f	an embedding
$\mathcal{N}_a$	neighborhood around node a in a graph
$R_{a,b}$	the smallest set of closest points to node a in an embedding f that contains node b
$\mathrm{MAP}(f)$	the mean average precision fidelity measure of the embedding f
$D(f)$	the distortion fidelity measure of the embedding f
$D_{\mathrm{wc}}(f)$	the worst-case distortion fidelity measure of the embedding f
G	a graph, typically with node set V and edge set E
T	a tree
a, b, c	nodes in a graph or tree
$\deg(a)$	the degree of node a
$\deg_{\mathrm{max}}$	maximum degree of a node in a graph
ℓ	the longest path length in a graph
τ	the scaling factor of an embedding
$\mathrm{reflect}_{x \to y}$	a reflection of x onto y in hyperbolic space
$\arg(z)$	the angle that the point z in the plane makes with the x-axis
X	matrix of points in hyperbolic space
Y	matrix of transformed distances
S	diagonal scaling matrix used in h-MDS
v	vector of squared norm values used in h-MDS
$u, \hat{u}$	eigenvectors used in h-MDS
Z	reduced matrix used in h-MDS
α, β	intermediate scalars used in h-MDS
γ	geodesic used in PGA
w_i	transformed points used in PGA

Table 5: Glossary of variables and symbols used in this paper.

B   Related Work

Our study of representation tradeoffs for hyperbolic embeddings was motivated by exciting recent approaches towards such embeddings in Nickel and Kiela [16] and Chamberlain et al. [4]. Earlier efforts proposed using hyperbolic spaces for routing, starting with Kleinberg's work on geographic routing [14]. Cvetkovski and Crovella [7] performed hyperbolic embeddings and routing for dynamic networks. Recognizing that the use of hyperbolic space for routing required a large number of bits to store the vertex coordinates, Eppstein and Goodrich [8] introduced a scheme for succinct embedding and routing in the hyperbolic plane.

Several papers have studied the notion of hyperbolicity of networks, starting with the seminal work on hyperbolic graphs Gromov [10]. More recently, Chen et al. [5] considered the hyperbolicity of small world graphs and tree-like random graphs. Abu-Ata and Dragan [1] performed a survey that examines how well real-world networks can be

approximated by trees using a variety of tree measures and tree embedding algorithms. To motivate their study of tree metrics, I. Abraham et al. [12] computed a measure of tree likeness on a Internet infrastructure network.

We use matrix completion to perform embeddings with incomplete data. Matrix completion is a celebrated problem. Candes and Tao [3] derive bounds on the minimum number of entries needed for completion for a fixed rank matrix; they also introduce a convex program for matrix completion operating at near the optimal rate.

Principal geodesic analysis (PGA) generalizes principal components analysis (PCA) for the manifold setting. It was introduced and applied to shape analysis in [9] and extended to a probabilistic setting in [21].

C Low-Level Formulation Details

We plan to release our PyTorch code, high precision solver, and other routines on Github. A few comments are helpful to understand the reformulation. In particular, we simply minimize the squared hyperbolic distance with a learned scale parameter, τ, e.g., :

$$\min_{x_1,\ldots,x_n,\tau} \sum_{1 \le i < j \le n} (\tau d_H(x_i, x_j) - d_{i,j})^2$$

We typically require that $\tau \ge 0.1$.

- On continuity of the derivative of the loss: Note that

$$\partial_x \mathsf{acosh}(1 + x) = \frac{1}{\sqrt{(1+x)^2 - 1}} = \frac{1}{\sqrt{x(x+2)}} \text{ hence } \lim_{x \to 0} \partial_x \mathsf{acosh}(1 + x) = \infty.$$

  Thus, $\lim_{y \to x} \partial_x d_H(x, y) = \infty$. In particular, if two points happen to get near to one another during execution, gradient-based optimization becomes unstable. Note that $\exp\{\mathsf{acosh}(1 + \mathrm{x})\}$ suffers from a similar issue, and is used in both [4, 16]. This change may increase numerical instability, and the public code for these approaches does indeed take steps like masking out updates to mitigate NANs. In contrast, the following may be more stable:

$$\partial_x \mathsf{acosh}(1 + x)^2 = 2 \frac{\mathsf{acosh}(1 + x)}{\sqrt{x(x+2)}} \text{ and in particular } \lim_{x \to 0} \partial_x \mathsf{acosh}(1 + x)^2 = 2$$

  The limits follows by simply applying L'Hopital's rule. In turn, this implies the square formulation is continuously differentiable. Note that it is not convex.

- One challenge is to make sure the gradient computed by PyTorch has the appropriate curvature correction (the Riemannian metric), as is well explained by Nickel and Kiela [16]. The modification is straightforward: we create a subclass of NN.PARAMETER called HYPERBOLIC_PARAMETER. This wrapper class allows us to walk the tree to apply the appropriate correction to the metric (which amounts to multiplying $\nabla_w f(w)$ by $\frac{1}{4}(1 - \|w\|^2)^2$). After calling the BACKWARD function, we call a routine to walk the autodiff tree to find such parameters and correct them. This allows HYPERBOLIC_PARAMETER and traditional parameters to be freely mixed.

- We project back on the hypercube following Nickel and Kiela [16] and use gradient clipping with bounds of $[-10^5, 10^5]$. This allows larger batch sizes to more fully utilize the GPU.

D Combinatorial Construction Proofs

We first derive a lower bound on the bits of precision required for embedding a graph into $\mathbb{H}_2$. Afterwards we prove a result bounding the precision for our extension of Sarkar's construction for the r-dimensional Poincaré ball $\mathbb{H}_r$. Finally, we give some details on the algorithm for this extension.

We derive the lower bound by exhibiting an explicit graph and lower bounding the precision needed to represent its nodes. The graph has a root node and 4 chains attached to the root, with n nodes each (for a total of $4n + 1$ nodes).
Lemma D.1. *The bits of precision needed to embed a graph with longest path ℓ is $\Omega\left(\frac{1}{\varepsilon}\ell\right)$.*

Proof. We start with the case $n = 1$, where we have the root and 4 leaves.

Note that all of the leaves should have equal norm $\|u\|$, or otherwise we could equalize this norm without increasing the distortion. To see this, let the longest edge have (in hyperbolic distance) length s and the shortest have length s'. The longest path has length at most $2s$, while the shortest path has length at least s'. The worst-case distortion is the largest expansion (at most $2s/2 = s/1$) multiplied by the largest contraction (at worst $2/(2s') = 1/s'$, or s/s'. Thus, equalizing s/s' can only decrease the distortion.

We may place the root at the origin. The largest separation we can retrieve between leaves is by aligning the chains with the axis. Now we are ready to examine distortion. Let x and y be a pair of closest leaves. We set $u = \|x\| = \|y\|$. Write $\bar{u} = \frac{u^2}{1-u^2}$. Then,

$$d_H(0, y) = d_H(0, x) = \operatorname{acosh}\left(1 + 2\bar{u}\right).$$

The distance between x and y is given by

$$
\begin{aligned}
d_H(x, y) &= \operatorname{acosh}\left(1 + 2\frac{\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)}\right) \\
&= \operatorname{acosh}\left(1 + 4(1 - \hat{x}^T\hat{y})\frac{u^2}{(1 - u^2)^2}\right),
\end{aligned}
$$

where $\hat{x}u = x$ and $\hat{y}u = y$.

Now, since $d_H(x, y) \le d_H(0, x) + d_H(0, y) = 2d_H(0, x)$, in order to achieve distortion ε we need

$$d_H(x, y) \ge 2d_H(0, x)(1 - \varepsilon).$$

Since $\cosh$ is monotonic, we can write

$$
\begin{aligned}
4\frac{u^2}{(1 - u^2)^2} &\ge 4(1 - \hat{x}^T\hat{y})\frac{u^2}{(1 - u^2)^2} \\
&\ge \cosh(2d_H(0, x)(1 - \varepsilon)) - 1 \\
&= 2\left(\cosh^2(d_H(0, x)(1 - \varepsilon)) - 1\right).
\end{aligned}
$$

Next, we use the bound

$$\cosh^2(z(1 - \varepsilon)) \ge (1 - 2\varepsilon\cosh(z))\cosh^2(z).$$

Applying this bound with $z = \operatorname{acosh}(1 + 2\bar{u})$, we have that

$$1 + 2\frac{\bar{u}^2}{u^2} \ge (1 - 2\varepsilon(1 + 2\bar{u}))(1 + 2\bar{u})^2. \tag{4}$$

Note that $\frac{1}{u^2} - 1 = \frac{1-u^2}{u^2} = \bar{u}^{-1}$. This yields

$$\frac{\bar{u}^2}{u^2} = \frac{\bar{u}^2}{u^2} - \bar{u}^2 + \bar{u}^2 = \bar{u}^2\left(\frac{1}{u^2} - 1\right) + \bar{u}^2 = \bar{u} + \bar{u}^2.$$

Applying this to (4),

$$1 + 2(\bar{u} + \bar{u}^2) \ge (1 - 2\varepsilon(1 + 2\bar{u}))(1 + 2\bar{u})^2.$$

Here the right-hand side can be written as $1 + 4\bar{u} + 4\bar{u}^2 - 2\varepsilon(1 + 2\bar{u})^3$. Then, simplifying, we have that

$$(\bar{u} + \bar{u}^2) \ge 2(\bar{u} + \bar{u}^2) - \varepsilon(1 + 2\bar{u})^3,$$

which can be rewritten as

$$\varepsilon(1 + 2\bar{u})^3 \ge (\bar{u} + \bar{u}^2).$$

Then,

$$8\bar{u} + o(\bar{u}) \geq \frac{1}{\epsilon}.$$

Using our earlier arguments about the bits of precision needed for an edge, $\bar{u} = \Omega(\epsilon^{-1})$ bits are required.

Next, we consider $n > 1$. Let us call one of our two chains $x_1, x_2, \ldots, x_n$ and another $y_1, y_2, \ldots, y_n$. We write $u = \|x_{n-1}\|$ and $v = \|x_n - x_{n-1}\|$, and set $\bar{u}$ and $\bar{v}$ as above. We proceed inductively. The $n = 1$ base case is complete. Assume $\bar{u} = \Omega((n-1)\varepsilon^{-1})$. Now consider $\bar{v}$. We can reflect x_{n-1} to the origin and follow the $n = 1$ argument above to recover that $\bar{v} = \Omega(\varepsilon^{-1})$. Now, since x_{n-1} and $x_n - x_{n-1}$ are collinear (as are y_n and $y_n - y_{n-1}$), we have that

$$
\begin{aligned}
\frac{\|x_n\|^2}{1 - \|x_n\|^2} &= \frac{\|x_{n-1} + (x_n - x_{n-1})\|^2}{1 - \|x_{n-1} + (x_n - x_{n-1})\|^2} \\
&\geq \frac{\|x_{n-1}\|^2}{1 - \|x_{n-1} + (x_n - x_{n-1})\|^2} + \frac{\|x_n - x_{n-1}\|^2}{1 - \|x_{n-1} + (x_n - x_{n-1})\|^2} \\
&\geq \frac{\|x_{n-1}\|^2}{1 - \|x_{n-1}\|^2} + \frac{\|x_n - x_{n-1}\|^2}{1 - \|x_n - x_{n-1}\|^2} \\
&= \bar{u} + \bar{v} = \Omega((n-1)\varepsilon^{-1}) + \Omega(\varepsilon^{-1}) = \Omega(n\varepsilon^{-1}).
\end{aligned}
$$

To complete the proof, note that the longest path is $\ell = 2n + 1$. $\qquad\square$

Next, we prove our extension of Sarkar's construction for $\mathbb{H}_r$, restated below.

Proposition 3.1. *The generalized $\mathbb{H}_r$ combinatorial construction has distortion at most $1 + \varepsilon$ and requires at most $O(\frac{1+\varepsilon}{\varepsilon} \frac{\ell}{r} \log \deg_{\max})$ bits to represent a node component for $r \leq (\log \deg_{\max}) + 1$, and $O(\frac{1+\varepsilon}{\varepsilon} \ell)$ bits for $r > (\log \deg_{\max}) + 1$.*

Proof. The combinatorial construction achieves worst-case distortion bounded by $1 + \varepsilon$ in two steps [18]. First, it is necessary to scale the embedded edges by a factor of τ sufficiently large to enable each child of a parent node to be placed in a disjoint cone. The smallest angle θ for a cone must be less than $\frac{\pi}{\deg_{\max}}$. The connection between this angle and the scaling factor τ relies on the relationship $\tau = -\log(\tan\theta/2)$. Note that the larger $\deg_{\max}$, the smaller θ, and the larger τ becomes.

This initial step provides a Delaunay embedding (and thus a MAP of 1.0), but perhaps not sufficient distortion. The second step is to further scale the points by a factor of $\frac{1+\varepsilon}{\varepsilon}$; this ensures the distortion upper bound.

Our generalization to the Poincaré ball of dimension r will modify the first step by showing that we can pack more children around a parent while maintaining the same angle. We use a generalization of cones for $\mathbb{H}_r$. We define cones $C(X, Y, \theta)$ for $\theta \in [0, \pi/2]$ as follows: $C(X, Y, \theta) = \left\{ Z \in \mathbb{H}_r : \langle Z - X, Y - X \rangle \geq \|Z - X\|\|Y - X\| \cos \frac{\theta}{2} \right\}$.

We use the following lower bound [13] on the maximum number of unit vectors $A(r, \theta)$ that can be placed on the unit sphere of dimension r with pairwise angle at least θ:

$$A(r, \theta) \geq (1 + o(1))\sqrt{2\pi r} \frac{\cos \theta}{(\sin \theta)^{r-1}}.$$

We consider angles $\theta \leq \pi/4$ and large r, so that

$$A(r, \theta) \geq \frac{1}{(\sin \theta)^{r-1}}.$$

We need to place as many as $\deg_{\max}$ vertices, so we require θ small enough so that

$$A(r, \theta) \geq \frac{1}{(\sin \theta)^{r-1}} \geq \deg_{\max}.$$

This condition is satisfied by

$$\theta = \operatorname{asin}(\deg_{\max}^{-\frac{1}{r-1}}).$$

17

Note the key difference compared to the two-dimensional case where $\theta = \frac{\pi}{\deg_{\max}}$; here we reduce the angle's dependence on the degree by an exponent of $\frac{1}{r-1}$. It remains to compute the explicit scaling factor τ that this angle yields; recall that τ satisfies $\tau = -\log(\tan\theta/2)$.

We then have

$$
\tau = -\log(\tan\theta/2) = -\log\left(\frac{\sin\theta}{1+\cos\theta}\right) = -\log\left(\frac{\sin\theta}{1+\sqrt{1-\sin^2\theta}}\right)
$$

$$
= -\log\left(\frac{1}{\deg_{\max}^{\frac{1}{r-1}} + \sqrt{\deg_{\max}^{\frac{2}{r-1}} - 1}}\right)
$$

$$
\leq \log\left(2\sqrt{\deg_{\max}^{\frac{2}{r-1}} - 1}\right) = \log 2 + O\left(\frac{1}{r}\log\deg_{\max}\right).
$$

This quantity tells us the scaling factor without considering distortion (the first step). To yield the $1+\varepsilon$ distortion, we must increase the scaling by a factor of $\frac{1+\varepsilon}{\varepsilon}$. The longest distance in the graph is the longest path ℓ multiplied by this quantity.

Putting it all together, for a tree with longest path ℓ, maximum degree $\deg_{\max}$ and distortion at most $1+\varepsilon$, the components of the embedding require (using the fact that distances $\|d\|$ require d bits),

$$
O\left(\frac{1+\varepsilon}{\varepsilon}\frac{\ell}{r}\log d_{\max}\right)
$$

bits per component for $r \leq (\log\deg_{\max}) + 1$. However, once we have increased the angles past $r = \log d_{\max}$ dimensions, the points cannot be further separated, and additional dimensions do not help. $\qquad\square$

Next, we provide more details on the coding-theoretic child placement construction for r-dimensional embeddings. Recall that children are placed at the vertices of a hypercube inscribed into the unit hypersphere, with components in $\frac{\pm 1}{\sqrt{r}}$. These points are indexed by sequences $\mathbf{a} \in \{0,1\}^r$ so that

$$
\mathbf{x_a} = \left(\frac{(-1)^{a_1}}{\sqrt{r}}, \frac{(-1)^{a_2}}{\sqrt{r}}, \ldots, \frac{(-1)^{a_r}}{\sqrt{r}}\right).
$$

The Euclidean distance between x_a and x_b is a function of the Hamming distance $d_{\text{Hamming}}(a,b)$ between a and b. The Euclidean distance is exactly $2\sqrt{\frac{d_{\text{Hamming}}(a,b)}{r}}$. Therefore, we can control the distances between the children by selecting a set of binary sequences with a prescribed minimum Hamming distance—a binary error-correcting code—and placing the children at the resulting hypercube vertices.

We introduce a small amount of terminology from coding theory. A binary code $\mathcal{C}$ is a set of sequences $\mathbf{a} \in \{0,1\}^r$. A $[r,k,h]_2$ code $\mathcal{C}$ is a binary linear code with length r (i.e., the sequences are of length r), size 2^k (there are 2^k sequences), and minimum Hamming distance h (the minimum Hamming distance between two distinct members of the code is h).

The Hadamard code $\mathcal{C}$ has parameters $[2^k, k, 2^{k-1}]$. If $r = 2^k$ is the dimension of the space, the Hamming distance between two members of $\mathcal{C}$ is at least $2^{k-1} = r/2$. Then, the distance between two distinct vertices of the hypercube x_a and x_b is $2\sqrt{\frac{r/2}{r}} = 2\sqrt{1/2} = \sqrt{2}$. Moreover, we can place up to $2^k = r$ points at least at this distance.

To build intuition, consider placing children on the unit circle ($r = 2$) compared to the $r = 128$-dimensional unit sphere. For $r = 2$, we can place up to 4 points with pairwise distance at least $\sqrt{2}$. However, for $r = 128$, we can place up to 128 children while maintaining this distance.

We briefly describe a few more practical details. Note that the Hadamard code is parametrized by k. To place $c+1$ children, take $k = \lceil\log_2(c+1)\rceil$. However, the desired dimension r' of the embedding might be larger than the resulting code length $r = 2^k$. We can deal with this by repeating the codeword. If there are r' dimensions and $r|r'$, then the

distance between the resulting vertices is still at least $\sqrt{2}$. Also, recall that when placing children, the parent node has already been placed. Therefore, we perform the placement using the hypercube, and rotate the hypersphere so that one of the $c + 1$ placed nodes is located at this parent.

E   Proof of Easy Recovery and h-MDS Results

In this section, we prove the results that we presented in Section 4.1. We will do this in two steps. First, we define the *easy recovery property* more formally, and restate and prove the result that every distance matrix that admits an exact embedding also admits an exact embedding that has the easy recovery property. Then, we derive Algorithm 2, h-MDS, and prove that it does produce an exact recovery whenever one is possible.

We start by stating the easy recovery property and proving the easy recovery lemma.

Definition E.1. *Let $X \in \mathbb{R}^{n \times r}$ be the matrix corresponding to n points $x_1, \ldots, x_n$ in the Poincare ball of dimension r. Given X, let $D \in \mathbb{R}^{n \times n}$ be the diagonal matrix defined as*

$$D_{i,i} = \frac{1}{1 - \|x_i\|^2},$$

let $v \in \mathbb{R}^n$ be the vector defined as

$$v_i = \|x_i\|^2,$$

and let u be the principle unit eigenvector of the matrix

$$R = D(1v^T + v1^T)D,$$

which is guaranteed by Perron-Frobenius to have nonnegative entries. We say that X has the easy-recovery property *if*

$$X^T D u = 0.$$

Lemma E.2. *If a distance matrix $d_{i,j}$ has an exact hyperbolic embedding, then it has an exact embedding with the easy-recovery property.*

Proof. Let $X \in \mathbb{R}^{n \times r}$ be the exact embedding, a matrix corresponding to to n points $x_1, \ldots, x_n$ in the Poincare ball of dimension r such that $d_H(x_i, x_j) = d_{i,j}$. Let $D(X) \in \mathbb{R}^{n \times n}$ be the diagonal matrix defined as

$$D_{i,i}(X) = \frac{1}{1 - \|x_i\|^2}$$

and let $v(X) \in \mathbb{R}^n$ be the vector defined as

$$v_i(X) = \|x_i\|^2.$$

Finally, let $u(X)$ be the principal unit eigenvector of the matrix

$$R(X) = D(X)(1v(X)^T + v(X)1^T)D(X),$$

which is guaranteed by Perron-Frobenius to have nonnegative entries. For any c on the Poincare ball let $T_c(w)$ denote the translation map in hyperbolic space that maps the point 0 to the point c. Let $y_i = T_c(x_i)$ for all $i \in \{1, \ldots, n\}$, and define matrix $Y \in \mathbb{R}^{n \times r}$ similarly to X. Note first that Y must also be an exact embedding of the distance matrix, since the translation mapping in hyperbolic space is isometric. If we could find a c such that

$$Y^T D(Y)u(Y) = 0$$

then this would prove our lemma statement. The remainder of the proof is to show that such a c exists.

First, extend T_c to matrix arguments so that $T_c(W)$ denotes the matrix formed by applying the translation operator T_c to all the vectors that make up W. (For example, in the lemma statement, $T_c(X) = Y$.) Next, define the map $M : \mathbb{R}^r \to \mathbb{R}^r$ as

$$M(c) = T_c(X)^T \frac{D(T_c(X))u(T_c(x))}{1^T D(T_c(X))u(T_c(x))}.$$

19

Our lemma statement is equivalent to saying that there exists a c such that $M(c) = 0$. Since both D and u are nonnegative by definition for points on the Poincare ball, it follows that

$$\rho(c) = \frac{D(T_c(X))u(T_c(x))}{1^T D(T_c(X))u(T_c(x))}$$

is a nonnegative vector that sums to 1. Now, M must be a continuous function for c on the Poincare ball, because it is composed of continuous functions (the only function that is not obviously continuous is u, but this can also be shown to be continuous using its explicit derivation). Furthermore, the range of M is also the Poincare ball, since $T_c(X)$ maps all the points of X to another point on the Poincare ball, and multiplying by ρ takes a convex combination of these points. So we've just shown that M is a continuous function from the open unit ball to the open unit ball.

Next, extend the definition of M to the boundary of the ball, by letting

$$M(w) = w$$

for any unit vector w. M will also be continuous (now in Euclidean space) at the boundary of the ball, because if we let c approach w, then all the points of $T_c(X)$ will be translated far in the w direction and so approach w in the Poincare ball. We can see this more explicitly by looking at the explicit formula for T_c in the Poincare ball,

$$T_c(x) = \frac{(1 + 2c^T x + \|x\|^2)c + (1 - |c|^2)x}{1 + 2c^T x + |c|^2 |x|^2},$$

which clearly approaches w as c approaches w for any x if $\|w\| = 1$.

We now have that $M(c)$ is a continuous mapping from the closed unit ball to the closed unit ball, and furthermore M is the identity function on the boundary of the ball, the unit sphere. Consider any continuous deformation of the unit sphere to a point on its interior (always remaining within the unit ball). The image of this deformation under M must *also* be a continuous deformation of the unit sphere to a point on its interior (always remaining within the unit ball). But such deformations must include all points in the unit ball, so it follows that the range of M is the entire unit ball. As a consequence 0 is in the range of M, so there exists a c such that $M(c) = 0$. This proves the lemma. $\square$

Now we show in greater detail how we can do exact recovery, by going through the derivation of h-MDS. Suppose that we are given distances $d_{i,j}$ that have an exact embedding as a set of hyperbolic points $x_1, \ldots, x_n \in \mathbb{H}_r$. Let $X \in \mathbb{R}^{n \times r}$ denote the matrix that corresponds to these points' mapping onto the Poincare ball model. Without loss of generality, suppose that X has the easy-recovery property: such an X is guaranteed to exist for any exactly-embeddable distance matrix by Lemma E.2. We shall see that this property implies that we can exactly recover x_i just from the observed distances. We start by defining the matrix $Y \in \mathbb{R}^{n \times n}$ where

$$Y_{i,j} = \frac{\cosh(d_{i,j}) - 1}{2} = \frac{\|x_i - x_j\|^2}{(1 - \|x_i\|^2)(1 - \|x_j\|^2)}.$$

It is clear that Y can be computed directly as a function of just the observed distances $d_{i,j}$.

Expanding the squared norm, we get a decomposition of Y showing that the rank of $Y_{i,j}$ is less than $r + 2$. Let D be the diagonal matrix such that $D_{i,i} = (1 - \|x_i\|^2)^{-1}$ and let $v_i = \|x_i\|^2$. Then we can express Y as

$$Y = D \left(v1^T + 1v^T - 2XX^T\right) D.$$

The inner term is a scaled euclidean distance matrix.

We observe Y and we'd like to recover X. To do so, we leverage the easy-recovery property. Let u be a principal eigenvector (guaranteed by Perron-Frobenius), with corresponding eigenvalue $\lambda > 0$, of the matrix

$$R = D \left(v1^T + 1v^T\right) D.$$

Since X has the easy-recovery property, $X^T D u = 0$, and so

$$Yu = D \left(v1^T + 1v^T - 2XX^T\right) D = Ru - 2DXX^T Du = Ru = \lambda u.$$

So u is also an eigenvector of Y with eigenvalue λ. Furthermore, it must be a principal eigenvector (again by Perron-Frobenius) since it has nonnegative entries. Since we can observe Y, we can also compute u directly via PCA. (This is the benefit of choosing X to have the easy-recovery property!) Suppose that we do this, and let $\hat{u}$ be the unit principle eigenvector returned by our procedure. Next, we want to recover D and v in terms of $\hat{u}$.

Since u was an eigenvector of R, and R is the symmetrization of a rank-1 matrix, we know that there must exist a principle eigenvector of the form $u = Dv + \alpha D1$ for some $\alpha > 0$. (This follows from the known properties of a symmetrization of a rank-1 matrix, and from the fact that Dv and $D1$ are both nonnegative.) Next, since $Dv = D1 - 1$, it follows that

$$Dv = \frac{u - \alpha 1}{1 + \alpha} \qquad \text{and} \qquad D1 = \frac{u + 1}{1 + \alpha}.$$

So, since u is an eigenvector of R with eigenvalue λ,

$$\lambda Dv + \lambda \alpha D1 = \lambda u = Ru = D1 \left(\frac{u - \alpha 1}{1 + \alpha} \right)^T u + Dv \left(\frac{u + 1}{1 + \alpha} \right)^T u.$$

Looking at the Dv and $D1$ components separately,

$$\lambda = \left(\frac{u + 1}{1 + \alpha} \right)^T u \qquad \text{and} \qquad \lambda \alpha = \left(\frac{u - \alpha 1}{1 + \alpha} \right)^T u.$$

Simplifying,

$$\lambda(1 + \alpha) = u^T u + 1^T u \qquad \text{and} \qquad \lambda \alpha (1 + \alpha) = u^T u - \alpha 1^T u.$$

Subtracting the second equation from the first produces

$$\lambda(1 - \alpha)(1 + \alpha) = (1 + \alpha) 1^T u$$

which simplifies to

$$\lambda(1 - \alpha) = 1^T u.$$

Subtracting this from the first component above produces

$$2\lambda \alpha = u^T u.$$

Thus,

$$\frac{(1^T u)^2}{u^T u} = \frac{\lambda^2 (1 - \alpha)^2}{2\lambda \alpha} = \frac{\lambda(1 - \alpha)^2}{2\alpha}$$

and so

$$\frac{(1 - \alpha)^2}{\alpha} = \frac{2(1^T u)^2}{\lambda u^T u}.$$

We can solve this directly using the quadratic formula. If we let

$$b = 1 + \frac{(1^T u)^2}{\lambda u^T u},$$

then regardless of what eigenvector we observe, we can compute b, because the expression for b is invariant to the scaling of u. Since u is just a scaling of our observed eigenvector $\hat{u}$, and this expression is a scale-invariant function of u, we can compute b directly as

$$b = 1 + \frac{(1^T \hat{u})^2}{\lambda \hat{u}^T \hat{u}},$$

In b, our equation becomes

$$(1 - \alpha)^2 = 2\alpha(b - 1) \; \to \; 0 = \alpha^2 - 2\alpha b + 1$$

which has two solutions, but the one in range is

$$\alpha = b - \sqrt{b^2 - 1}.$$

This is the only one in range because $b > 1$ and we need $\alpha \le 1$ for the sign to check out on $\lambda(1-\alpha) = 1^T u$. Once we have alpha, we can re-scale the eigenvector we observe to compute u directly, using

$$u = \frac{\hat{u}}{1^T \hat{u}} \lambda(1-\alpha),$$

and from there we can find Dv, $D1$, and v using the expressions above. Finally, we can compute XX^T as

$$XX^T = \frac{v1^T + 1v^T - D^{-1}YD^{-1}}{2},$$

and using PCA again we can compute X (or at least an isometric rotation of X) in terms of XX^T. This X will be an exact embedding of our original distance matrix d, and this X is exactly the value that is output by h-MDS.

F Perturbation Analysis

F.1 Handling Perturbations

Now that we have shown that h-MDS recovers an embedding exactly, we consider the impact of perturbations on the data. Given the necessity of high precision for some embeddings, we expect that in some regimes the algorithm should be very sensitive. Our results identify the scaling of those perturbations.

First, we consider how to measure the effect of a perturbation on the resulting embedding. We measure the gap between two configurations of points, written as matrices in $\mathbb{R}^{n \times r}$, by the sum of squared differences $D(X, Y) = \text{trace}((X-Y)^T(X-Y))$. Of course, this is not immediately useful, since X and Y can be rotated or reflected without affecting the distance matrix used for MDS–as these are isometries, while scalings and Euclidean translations are not. Instead, we measure the gap by
$$D_E(X, Y) = \inf\{D(X, PY) : P^T P = I\}.$$

In other words, we look for the configuration of Y with the smallest gap relative to X. For Euclidean MDS, Sibson [19] provides an explicit formula for $D_E(X, Y)$ and uses this formulation to build a perturbation analysis for the case where Y is a configuration recovered by performing MDS on the perturbed matrix $XX^T + \Delta(E)$, with $\Delta(E)$ symmetric.

Problem setup. In our case, the perturbations affect the hyperbolic distances. Let $H \in \mathbb{R}^{n \times n}$ be the distance matrix for a set of points in hyperbolic space. Let $\Delta(H) \in \mathbb{R}^{n \times n}$ be the perturbation, with $H_{i,i} = 0$ and $\Delta(H)$ symmetric (so that $\hat{H} = H + \Delta_H$ remains symmetric). To simplify our derivations, we assume that the perturbation of H does not alter the dominant (Perron-Frobenius) eigenvalue and eigenvector of Y from (1). The goal of our analysis is to estimate the gap $D_E(X, Y)$ between X recovered from H with h-MDS and $\hat{X}$ recovered from the perturbed distances $H + \Delta(H)$.

Lemma F.1. *Under the above conditions, if $\lambda_{\min}$ denotes the smallest nonzero eigenvalue of XX^T then up to second order in $\Delta(H)$,*

$$D_E(X, \hat{X}) \le \frac{2n^2}{\lambda_{\min}} \sinh^2\left(\|H\|_\infty\right) \|\Delta(H)\|_\infty^2.$$

The proof of this Lemma is contained in the Appendix. However, the key takeaway is that this upperbound matches our intuition for the scaling: if all points are close to one another, then $\|H\|_\infty$ is small and the space is approximately flat (since $\sinh^2(z)$ is dominated by $2z^2$ close to the origin). While points at great distance are sensitive to perturbations in an absolute sense. Note that far away points may be represented by have very large norms in some choices of coordinates. This leads us to consider other notions of recovery, which we do next.

F.2 Proof of Lemma F.1

Proof. Similarly to our development of h-MDS, we proceed by accessing the underlying Euclidean distance matrix, and then apply the perturbation analysis from Sibson [20]. There are two steps: first, we get rid of the acosh in the distances to leave us with scaled Euclidean distances. Next, we remove the scaling factors, and apply Sibson's result.

Hyperbolic to scaled Euclidean distortion. Let Y denote the scaled-Euclidean distance matrix, as in (1), so that $Y_{i,j} = \frac{1}{2}(\cosh(H_{i,j}) - 1)$. Let $\hat{Y}_{i,j} = \frac{1}{2}(\cosh(H_{i,j} + \Delta(H)_{i,j}) - 1)$ We write $\Delta(Y) = \hat{Y} - Y$ for the scaled Euclidean version of the perturbation. We can use the hyperbolic-cosine difference formula on each term to write

$$\Delta(Y)_{i,j} = \frac{1}{2}(\cosh(\hat{H}_{i,j}) - 1) - \frac{1}{2}(\cosh(H_{i,j}) - 1)$$

$$= \frac{1}{2}(\cosh(H_{i,j} + \Delta(H)_{i,j}) - \cosh(H_{i,j}))$$

$$= \sinh\left(\frac{2H_{i,j} + \Delta(H)_{i,j}}{2}\right) \sinh\left(\frac{\Delta(H)_{i,j}}{2}\right).$$

in terms of the infinity norm, as long as $\|H\|_\infty \geq \|\Delta(H)\|_\infty$ (it is fine to assume this because we are only deriving a bound up to second order, so we can suppose that $\Delta(H)$ is small), we can simplify this to

$$\|\Delta(Y)\|_\infty \leq \sinh(\|H\|_\infty) \sinh(\|\Delta(H)\|_\infty).$$

Scaled Euclidean to Euclidean inner product. Next we want the errors $\Delta(G)$ in the underlying Euclidean setting. From (2) we have that if $G = XX^T$, then $Y = R - 2DGD$, for matrix D with diagonal entries $D_{i,i} = (1 - \|x_i\|^2)^{-1}$ and matrix R as defined in the easy recovery section. By assumption, since the dominant eigenvalue and eigenvector of Y are not affected by the perturbation, and D and R are functions of these, D and R will also be unaffected by the perturbation, so also $\hat{Y} = R - 2D\hat{G}D$, where $\hat{G} = \hat{X}\hat{X}^T$. Since all the x_i are in the Poincar'e disk, it follows that $\|x_i\|^2 \leq 1$, and so $D_{i,i}^{-1} = 1 - \|x_i\|^2 \in [0,1]$. As a result, if we define $\Delta(G) = \hat{G} - G$, then

$$\|\Delta(G)\|_\infty = \|2D^{-1}\Delta(Y)D^{-1}\|_\infty \leq 2\|\Delta(Y)\|_\infty$$

and so

$$\|\Delta(G)\|_\infty \leq 2\sinh(\|H\|_\infty) \sinh(\|\Delta(H)\|_\infty).$$

Now we are in the Euclidean setting, and we can thus measure the result of the perturbation on X. We can apply Theorem 4.1 in Sibson [20]. This result states that if the perturbation takes the inner products XX^T to $XX^T + \Delta(G)$, and $\hat{X}$ is the configuration recovered from the perturbed inner products, then, the lowest-order term in the expansion of the error $D_E(X, \hat{X})$ in the perturbation $\Delta(G)$ is

$$D_E(X, \hat{X}) = \frac{1}{2} \sum_{j,k} \frac{(v_j^T \Delta(G) v_k)^2}{\lambda_j + \lambda_k}.$$

Here, the λ_i and v_i are the eigenvalues and corresponding orthonormal eigenvectors of XX^T and the sum is taken over pairs of λ_j, λ_k that are not both 0. Let $\lambda_{\min}$ be the smallest nonzero eigenvalue of XX^T. Then,

$$D_E(X, \hat{X}) \leq \frac{1}{2\lambda_{\min}} \sum_{j,k} (v_j^T \Delta(G) v_k)^2 \leq \frac{1}{2\lambda_{\min}} \|\Delta(G)\|_F^2$$

$$\leq \frac{n^2}{2\lambda_{\min}} \|\Delta(G)\|_\infty^2.$$

Combining this with the previous bounds, and restricting to second-order terms in $\|\Delta(H)\|_\infty^2$ proves the lemma. $\qquad\square$

G   Proof of Lemma 4.1

In this section, we prove Lemma 4.1, which gives a setting under which we can guarantee that the hyperbolic PGA objective is locally convex.

Proof of Lemma 4.1. We begin by considering the component function

$$f_i(\gamma) = \operatorname{acosh}^2(1 + d_E^2(\gamma, v_i)).$$

Here, the γ is a geodesic through the origin. We can identify this geodesic on the Poincaré disk with a unit vector u such that $\gamma(t) = (2t - 1)u$. In this case, simple Euclidean projection gives us

$$d_E^2(\gamma, v_i) = \|(I - uu^T)v_i\|^2.$$

Optimizing over γ is equivalent to optimizing over u, and so

$$f_i(u) = \operatorname{acosh}^2\left(1 + \|(I - uu^T)v_i\|^2\right).$$

If we define the functions

$$h(\gamma) = \operatorname{acosh}^2(1 + \gamma)$$

and

$$R(u) = \|(I - uu^T)v_i\|^2 = \|v_i\|^2 - (u^T v_i)^2$$

then we can rewrite f_i as

$$f_i(u) = h(R(u)).$$

Now, optimizing over u is an geodesic optimization problem on the hypersphere. Every goedesic on the hypersphere can be isometrically parameterized in terms of an angle θ as

$$u(\theta) = x\cos(\theta) + y\sin(\theta)$$

for orthogonal unit vectors x and y. Without loss of generality, suppose that $y^T v_i = 0$ (we can always choose such a y because there will always be some point on the geodesic that is orthogonal to v_i). Then, we can write

$$R(\theta) = \|v_i\|^2 - (x^T v_i)^2 \cos^2(\theta) = \|v_i\|^2 - (x^T v_i)^2 + (x^T v_i)^2 \sin^2(\theta).$$

Differentiating the objective with respect to θ,

$$\frac{d}{d\theta} h(R(\theta)) = h'(R(\theta))R'(\theta)$$
$$= 2h'(R(\theta)) \cdot (v_i^T x)^2 \cdot \sin(\theta)\cos(\theta).$$

Differentiating again,

$$\frac{d^2}{d\theta^2} h(R(\theta)) = 4h''(R(\theta)) \cdot (v_i^T x)^4 \cdot \sin^2(\theta)\cos^2(\theta) + 2h'(R(\theta)) \cdot (v_i^T x)^2 \cdot \left(\cos^2(\theta) - \sin^2(\theta)\right).$$

Now, suppose that we are interested in the Hessian at a point $z = x\cos(\theta) + y\sin(\theta)$ for some fixed angle θ. Here, $R(\theta) = R(z)$, and as always $v_i^T z = v_i^T x\cos(\theta)$, so

$$\frac{d^2}{d\theta^2} h(R(\theta))\Big|_{u(\theta)=z} = 4h''(R(\theta)) \cdot (v_i^T x)^4 \cdot \sin^2(\theta)\cos^2(\theta) + 2h'(R(\theta)) \cdot (v_i^T x)^2 \cdot \left(\cos^2(\theta) - \sin^2(\theta)\right)$$
$$= 4h''(R(z)) \cdot \frac{(v_i^T z)^4}{\cos^4(\theta)} \cdot \sin^2(\theta)\cos^2(\theta) + 2h'(R(z)) \cdot \frac{(v_i^T x)^2}{\cos^2(\theta)} \cdot \left(\cos^2(\theta) - \sin^2(\theta)\right)$$
$$= 4h''(R(z)) \cdot (v_i^T z)^4 \cdot \tan^2(\theta) + 2h'(R(z)) \cdot (v_i^T z)^2 \cdot \left(1 - \tan^2(\theta)\right)$$
$$= 2h'(R(z)) \cdot (v_i^T z)^2 + \left(4h''(R(z)) \cdot (v_i^T z)^4 - 2h'(R(z)) \cdot (v_i^T z)^2\right) \tan^2(\theta).$$

But we know that since h is concave and increasing, this last expression in parenthesis must be negative. It follows that a lower bound on this expression for fixed z will be attained when $\tan^2(\theta)$ is maximized. For any geodesic through z, the angle θ is the distance along the geodesic to the point that is (angularly) closest to v_i. By the Triangle inequality, this will be no greater than the distance θ along the Geodesic that connects z with the normalization of v_i. On this worst-case geodesic,

$$v_i^T z = \|v_i\| \cos(\theta),$$

and so
$$\cos^2(\theta) = \frac{(v_i^T z)^2}{\|v_i\|^2}$$
and
$$\tan^2(\theta) = \sec^2(\theta) - 1 = \frac{\|v_i\|^2}{(v_i^T z)^2} - 1 = \frac{R(z)}{(v_i^T z)^2}.$$

Thus, for any geodesic, for the worst-case angle θ,

$$\frac{d^2}{d\theta^2} h(R(\theta))\big|_{u(\theta)=z} \geq 2h'(R(z)) \cdot (v_i^T z)^2 + \left(4h''(R(z)) \cdot (v_i^T z)^4 - 2h'(R(z)) \cdot (v_i^T z)^2\right)\tan^2(\theta)$$
$$= 2h'(R(z)) \cdot (v_i^T z)^2 + \left(4h''(R(z)) \cdot (v_i^T z)^2 - 2h'(R(z))\right) R(z).$$

From here, it is clear that this lower bound on the second derivative (and as a consequence local convexity) is a function solely of the norm of v_i and the residual to z. From simple evaluation, we can compute that

$$h'(\gamma) = 2\frac{\mathrm{acosh}(1+\gamma)}{\sqrt{\gamma^2+2\gamma}}$$

and

$$h''(x) = 2\frac{\sqrt{\gamma^2+2\gamma} - (1+\gamma)\,\mathrm{acosh}(1+\gamma)}{(\gamma^2+2\gamma)^{3/2}}.$$

As a result

$$4\gamma h''(\gamma) + h'(\gamma) = 8\frac{\gamma\sqrt{\gamma^2+2\gamma} - (\gamma^2+\gamma)\,\mathrm{acosh}(1+\gamma)}{(\gamma^2+2\gamma)^{3/2}} + 2\frac{(\gamma^2+2\gamma)\,\mathrm{acosh}(1+\gamma)}{(\gamma^2+2\gamma)^{3/2}}$$
$$= 2\frac{4\gamma\sqrt{\gamma^2+2\gamma} - 4(\gamma^2+\gamma)\,\mathrm{acosh}(1+\gamma) + (\gamma^2+2\gamma)\,\mathrm{acosh}(1+\gamma)}{(\gamma^2+2\gamma)^{3/2}}$$
$$= 2\frac{4\gamma\sqrt{\gamma^2+2\gamma} - (3\gamma^2+2\gamma)\,\mathrm{acosh}(1+\gamma)}{(\gamma^2+2\gamma)^{3/2}}.$$

For any γ that satisfies $0 \leq \gamma \leq 1$,

$$4\gamma\sqrt{\gamma^2+2\gamma} \geq (3\gamma^2+2\gamma)\,\mathrm{acosh}(1+\gamma)$$

and so

$$4\gamma h''(\gamma) + h'(\gamma) \geq 0.$$

Thus, if $0 \leq R(z) \leq 1$,

$$\frac{d^2}{d\theta^2} h(R(\theta))\big|_{u(\theta)=z} \geq 2h'(R(z)) \cdot (v_i^T z)^2 + \left(4h''(R(z)) \cdot (v_i^T z)^2 - 2h'(R(z))\right) R(z)$$
$$= h'(R(z)) \cdot (v_i^T z)^2 + \left(4h''(R(z)) \cdot R(z) + h'(R(z))\right) \cdot (v_i^T z)^2 - 2h'(R(z)) \cdot R(z)$$
$$\geq h'(R(z)) \cdot (v_i^T z)^2 - 2h'(R(z)) \cdot R(z)$$
$$= h'(R(z)) \cdot \left(\|v_i\|^2 - R(z)\right) - 2h'(R(z)) \cdot R(z)$$
$$= h'(R(z)) \cdot \left(\|v_i\|^2 - 3R(z)\right).$$

Thus, a sufficient condition for convexity is for (as we assumed above) $R(z) \leq 1$ and

$$\|v_i\|^2 \geq 3R(z).$$

Combining these together shows that if

$$\mathrm{acosh}^2\left(1 + d_E(\gamma, v_i)^2\right) = R(z) \leq \min\left(1, \frac{1}{3}\|v_i\|^2\right)$$

then f_i is locally convex at z. The result of the lemma now follows from the fact that f is the sum of many f_i and the sum of convex functions is also convex. $\qquad\square$

H Experimental Results

Combinatorial Construction: Parameters To improve the intuition behind the combinatorial construction, we report some additional parameters used by the construction. For each of the graphs, we report the maximum degree, the scaling factor ν that the construction used (note how these vary with the size of the graph and the maximal degree), the time it took to perform the embedding, in seconds, and the number of bits needed to store a component for $\varepsilon = 0.1$ and $\varepsilon = 1.0$.

Dataset	Nodes	Edges	$d_{\max}$	Time	$\varepsilon = 0.1$ Scaling Factor	Precision	$\varepsilon = 1.0$ Scaling Factor	Precision
Bal. Tree 1	40	39	4	3.78	23.76	102	4.32	18
Phylo. Tree	344	343	16	3.13	55.02	2361	10.00	412
WordNet	74374	75834	404	1346.62	126.11	2877	22.92	495
CS PhDs	1025	1043	46	4.99	78.30	2358	14.2	342
Diseases	516	1188	24	3.92	63.97	919	13.67	247
Protein - Yeast	1458	1948	54	6.23	81.83	1413	15.02	273
Gr-QC	4158	13428	68	75.41	86.90	1249	16.14	269
California	5925	15770	105	114.41	96.46	1386	19.22	245

Table 6: Combinatorial construction parameters and results.

We also comment on the possibility of reducing the precision by increasing the dimension (using the extended $\mathbb{H}_r$ construction from Section 3). We noted the largest dimension that reduces the precision (beyond this additional dimensions do not help). We see that the $\varepsilon = 0.1$ remains challenging to store, while with $\varepsilon = 1$, 64 bits suffices.

Dataset	Nodes	Edges	$d_{\max}$	Dimension	$\varepsilon = 0.1$ Precision	$\varepsilon = 1.0$ Precision
Bal. Tree 1	40	39	4	3	26	5
Phylo. Tree	344	343	16	5	251	46
WordNet	74374	75834	404	10	133	25
CS PhDs	1025	1043	46	7	176	32
Diseases	516	1188	24	6	85	16
Protein - Yeast	1458	1948	54	7	101	19
Gr-QC	4158	13428	68	8	86	16
California	5925	15770	105	8	84	16

Table 7: Reducing the precision with larger dimensions.

Hyperparameter: Effect of Rank We also considered the influence of the dimension on the perfomance of h-MDS, PCA, and FB. On the Phylogenetic tree dataset, we measured distortion and MAP metrics for dimensions of 2,5,10,50,100, and 200. The results are shown in Table 8. We expected all of the techniques to improve with better rank, and this was the case as well. Here, the optimization-based approach typically produces the best MAP, optimizing the fine details accurately. We observe that the gap is closed when considering 2-MAP (that is, MAP where the retrieved neighbors are at distance up to 2 away). In particular we see that the main limitation of h-MDS is at the finest layer, confirming the idea MAP is heavily influenced by local changes. In terms of distortion, we found that h-MDS offers good performance even at a very low dimension (0.083 at 5 dimensions).

Precision Experiment (cf Table 9).

Rank	MAP			2-MAP			d_{avg}		
	h-MDS	PCA	FB	h-MDS	PCA	FB	h-MDS	PCA	FB
Rank 2	0.346	0.614	**0.718**	0.754	**0.874**	0.802	**0.317**	0.888	0.575
Rank 5	0.439	0.627	**0.761**	0.844	0.905	**0.950**	**0.083**	0.833	0.583
Rank 10	0.471	0.632	**0.777**	0.857	0.912	**0.953**	**0.048**	0.804	0.586
Rank 50	0.560	0.687	**0.784**	0.880	0.962	**0.974**	**0.036**	0.768	0.584
Rank 100	0.645	0.698	**0.795**	0.926	**0.999**	0.981	**0.036**	0.760	0.583
Rank 200	0.823	**1.0**	0.811	0.968	**1.0**	0.986	**0.039**	0.746	0.583

Table 8: Phylogenetic tree dataset. Variation with rank, measured with MAP, 2-MAP, and d_{avg}.

Precision	D_{avg}	MAP
128	0.357	0.347
256	0.091	0.986
512	0.076	1.0
1024	0.064	1.0

Table 9: h-MDS recovery at different precision levels for a 3-ary tree and rank 10.