Hyperbolic Distance Matrices

Puoya Tabaghi

University of Illinois at Urbana-Champaign
tabaghi2@illinois.edu

Ivan Dokmanić

University of Basel
ivan.dokmanic@unibas.ch

ABSTRACT

Hyperbolic space is a natural setting for mining and visualizing data with hierarchical structure. In order to compute a hyperbolic embedding from comparison or similarity information, one has to solve a hyperbolic distance geometry problem. In this paper, we propose a unified framework to compute hyperbolic embeddings from an arbitrary mix of noisy metric and non-metric data. Our algorithms are based on semidefinite programming and the notion of a hyperbolic distance matrix, in many ways parallel to its famous Euclidean counterpart. A central ingredient we put forward is a semidefinite characterization of the hyperbolic Gramian—a matrix of Lorentzian inner products. This characterization allows us to formulate a semidefinite relaxation to efficiently compute hyperbolic embeddings in two stages: first, we complete and denoise the observed hyperbolic distance matrix; second, we propose a spectral factorization method to estimate the embedded points from the hyperbolic distance matrix. We show through numerical experiments how the flexibility to mix metric and non-metric constraints allows us to efficiently compute embeddings from arbitrary data.

CCS CONCEPTS

• **Human-centered computing** → **Hyperbolic trees**; • **Computing methodologies** → **Machine learning**; • **Networks**;

KEYWORDS

distance geometry, hyperbolic space, semidefinite program, spectral factorization

ACM Reference Format:

Puoya Tabaghi and Ivan Dokmanić. 2020. Hyperbolic Distance Matrices. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), August 23–27, 2020, Virtual Event, CA, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3394486.3403224>

1 INTRODUCTION

Hyperbolic space is roomy. It can embed hierarchical structures uniformly and with arbitrarily low distortion [30, 45]. Euclidean space cannot achieve comparably low distortion even using an unbounded number of dimensions [33].

Embedding objects in hyperbolic spaces has found a myriad applications in exploratory science, from visualizing hierarchical

structures such as social networks and link prediction for symbolic data [40, 53] to natural language processing [14, 31], brain networks [8], gene ontologies [4] and recommender systems [10, 54].

Commonly in these applications, there is a tree-like data structure which encodes *similarity* between a number of entities. We experimentally observe some relational information about the structure and the data mining task is to find a geometric representation of the entities consistent with the experimental information. In other words, the task is to compute an embedding. This concept is closely related to the classical distance geometry problems and multidimensional scaling (MDS) [29] in Euclidean spaces [16, 32].

The observations can be metric or non-metric. Metric observations convey (inexact) distances; for example, in internet distance embedding a small subset of nodes with complete distance information are used to estimate the remaining distances [47]. Non-metric observations tell us which pairs of entities are closer and which are further apart. The measure of closeness is typically derived from domain knowledge; for example, word embedding algorithms aim to relate semantically close words and their topics [36, 43].

In scientific applications it is desirable to compute good low-dimensional hyperbolic embeddings. Insisting on low dimension not only facilitates visualization, but also promotes simple explanations of the phenomenon under study. However, in most works that leverage hyperbolic geometry the embedding technique is not the primary focus and the related computations are often ad hoc. The situation is different in the Euclidean case, where the notions of MDS, Euclidean distance matrices (EDMs) and their characterization in terms of positive semidefinite Gram matrices play a central role in the design and analysis of algorithms [3, 32].

In this paper, we focus on computing low-dimensional hyperbolic embeddings. While there exists a strong link between Euclidean geometry and positive (semi)definiteness, we prove that what we call *hyperbolic distance matrices* (HDMs) can also be characterized via semidefinite constraints. Unlike in the Euclidean case, the hyperbolic analogy of the Euclidean Gram matrix is a linear combination of two rank-constrained semidefinite variables. Together with a spectral factorization method to directly estimate the hyperbolic points, this characterization gives rise to flexible embedding algorithms which can handle diverse constraints and mix metric and non-metric data.

1.1 Related Work

The usefulness of hyperbolic space stems from its ability to efficiently represent the geometry of complex networks [5, 28]. Embedding metric graphs with underlying hyperbolic geometry has applications in word embedding [36, 43], geographic routing [27], routing in dynamical graphs [12], odor embedding [61], internet network embedding for delay estimation and server selection [7, 47], to name a few. In the literature such problems are known as hyperbolic multidimensional scaling [13].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7998-4/20/08...\$15.00

<https://doi.org/10.1145/3394486.3403224>

There exist Riemann gradient-based approaches [11, 31, 40, 41] which can be used to directly estimate such embeddings from metric measurements [44]. We emphasize that these methods are iterative and only guaranteed to return a locally optimal solution. On the other hand, there exist one-shot methods to estimate hyperbolic embeddings from a *complete* set of measured distances. The method of Wilson *et al.* [55] is based on spectral factorization of an inner product matrix (we refer to it as hyperbolic Gramian) that directly minimizes a suitable *stress*. In this paper, we derive a semidefinite relaxation to estimate the *missing* measurements and denoise the distance matrix, and then follow it with the spectral embedding algorithm.

Non-metric (or order) embedding has been proposed to learn visual-semantic hierarchies from ordered input pairs by embedding symbolic objects into a low-dimensional space [52]. In the Euclidean case, stochastic triplet embeddings [50], crowd kernels [49], and generalized non-metric MDS [1] are some well-known order embedding algorithms. For embedding hierarchical structures, Ganea *et al.* [20] model order relations as a family of nested geodesically convex cones. Zhou *et al.* [61] show that odors can be efficiently embedded in hyperbolic space provided that the similarity between odors is based on the statistics of their co-occurrences within natural mixtures.

1.2 Contributions

We summarize our main contributions as follows:

- **Semidefinite characterization of HDMs:** We introduce HDMs as an elegant tool to formalize distance problems in hyperbolic space; this is analogous to Euclidean distance matrices (EDM). We derive a semidefinite characterization of HDMs by studying the properties of hyperbolic Gram matrices—matrices of Lorentzian (indefinite) inner products of points in a hyperbolic space.
- **A flexible algorithm for hyperbolic distance geometry problems (HDGPs):** We use the semidefinite characterization to propose a flexible embedding algorithm based on semidefinite programming. It allows us to seamlessly combine metric and non-metric problems in one framework and to handle a diverse set of constraints. The non-metric and metric measurements are imputed as linear and quadratic constraints.
- **Spectral factorization and projection:** We compute the final hyperbolic embeddings with a simple, closed-form spectral factorization method.¹ We also propose a suboptimal method to find a low-rank approximation of the hyperbolic Gramian in the desired dimension.

1.3 Paper Organization

We first briefly review the analytical models of hyperbolic space and formalize hyperbolic distance geometry problems (HDGPs) in Section 2. Our framework is parallel with semidefinite approaches for Euclidean distance problems as per Table 1. In the 'Loid model,

¹After posting the first version of our manuscript we became aware that such a one-shot spectral factorization technique was proposed at least as early as in [55]. The same technique is also used by [26].

Table 1: Essential elements in semidefinite approach for distance problems, Euclidean versus hyperbolic space.

Euclidean	Hyperbolic
Euclidean Distance Matrix Gramian	Hyperbolic Distance Matrix H-Gramian
Semidefinite relaxation to complete an EDM	Semidefinite relaxation to complete an HDM
Spectral factorization of a Gramian to estimate the points	Spectral factorization of an H-Gramian to estimate the points

we define hyperbolic distance matrices to compactly encode hyperbolic distance measurements. We show that an HDM can be characterized in terms of the matrix of *indefinite* inner products, the hyperbolic Gramian. In Section 3, we propose a semidefinite representation of hyperbolic Gramians, and in turn HDMs. We cast HDGPs as rank-constrained semidefinite programs, which are then convexified by relaxing the rank constraints. We develop a spectral method to find a sub-optimal low-rank approximation of the hyperbolic Gramian, to the correct embedding dimension. Lastly, we propose a closed-form factorization method to estimate the embedded points. This framework lets us tackle a variety of embedding problems, as shown in Section 4, with real (odors) and synthetic (random trees) data. The proofs of propositions and derivations of proposed algorithms are given in the appendix and a summary of used notations is given in Table 2.

Table 2: Summary of notations.

Symbol	Meaning
$[m]$	Short for $\{1, \dots, m\}$
$[M]_{as}^2$	Asymmetric pairs $\{(m, n) : m < n, m, n \in [M]\}$
$x = [x_0, \dots, x_{m-1}]^T$	A vector in $\mathbb{R}^m$
$X = (x_{i,j})_{i \in [m], j \in [n]}$	A matrix in $\mathbb{R}^{m \times n}$
$X \geq 0$	A positive semidefinite (square) matrix
$\ X\ _F$	Frobenius norm of X
$\ X\ _2$	Operator norm of X
$\ X\ _{1,2}$	The ℓ_2 norm of columns' ℓ_1 norms, $\ [\ x_1\ _1, \dots, \ x_n\ _1]^T\ _2$
$\mathbb{E}_N[x]$	Empirical expectation of a random variable, $N^{-1} \sum_{n=1}^N x_n$
$e_m \in \mathbb{R}^M$	The m -th standard basis vector in $\mathbb{R}^M$
$P_r(X)$	The projection of $X \geq 0$ onto the span of its top r eigenvectors
$\mathbf{1}$	All-one vector of appropriate dimension
$\mathbf{0}$	All-zero vector of appropriate dimension

2 HYPERBOLIC DISTANCE GEOMETRY PROBLEMS

2.1 Hyperbolic Space

Hyperbolic space is a simply connected Riemannian manifold with constant negative curvature [6, 9]. In comparison, Euclidean and elliptic geometries are spaces with zero (flat) and constant positive curvatures. There are five isometric models for hyperbolic space: half-space ($\mathbb{H}^d$), Poincaré (interior of the disk) ($\mathbb{I}^d$), hemisphere ($\mathbb{J}^d$), Klein ($\mathbb{K}^d$), and 'Loid ($\mathbb{L}^d$) [9] (Figure 1). Each provides unique insights into the properties of hyperbolic geometry.

In the machine learning community the most popular models of hyperbolic geometry are Poincaré and 'Loid. We work in the 'Loid model as it has a simple, tractable distance function. It lets us cast the HDGP (formally defined in Section 2.2) as a rank-constrained

Figure 1: Models of hyperbolic space with level sets (colors) illustrating isometries.

semidefinite program. Importantly, it also leads to a closed-form embedding by a spectral method. For better visualization, however, we map the final embedded points to the Poincaré model via the stereographic projection, see Sections 2.1.2 and 4.

2.1.1 'Loid Model. Let x and y be vectors in $\mathbb{R}^{d+1}$ with $d \geq 1$. The Lorentzian inner product of x and y is defined as

$$[x, y] = x^T H y, \quad (1)$$

where

$$H = \begin{pmatrix} -1 & 0^T \\ 0 & I \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}. \quad (2)$$

This is an indefinite inner product on $\mathbb{R}^{d+1}$. The Lorentzian inner product has almost all the properties of ordinary inner products, except that

$$\|x\|_H^2 \stackrel{\text{def}}{=} [x, x]$$

can be positive, zero, or negative. The vector space $\mathbb{R}^{d+1}$ equipped with the Lorentzian inner product (1) is called a Lorentzian $(d+1)$ -space, and is denoted by $\mathbb{R}^{1,d}$. In a Lorentzian space we can define notions similar to the Gram matrix, adjoint, and unitary matrices known from Euclidean spaces as follows.

DEFINITION 1 (H-ADJOINT [23]). The H -adjoint $R^{[*]}$ of an arbitrary matrix $R \in \mathbb{R}^{(d+1) \times (d+1)}$ is characterized by

$$[R x, y] = [x, R^{[*]} y], \quad \forall x, y \in \mathbb{R}^{d+1}.$$

Equivalently,

$$R^{[*]} = H^{-1} R^T H. \quad (3)$$

DEFINITION 2 (H-UNITARY MATRIX [23]). An invertible matrix R is called H -unitary if $R^{[*]} = R^{-1}$.

The 'Loid model of d -dimensional hyperbolic space is a Riemannian manifold $\mathcal{L}^d = (\mathbb{L}^d, (g_x)_x)$, where

$$\mathbb{L}^d = \{x \in \mathbb{R}^{d+1} : \|x\|_H^2 = -1, x_0 > 0\}$$

and $g_x = H$ is the Riemannian metric.

DEFINITION 3 (LORENTZ GRAMIAN, H-GRAMIAN). Let the columns of $X = [x_1, x_2, \dots, x_N]$ be the positions of N points in $\mathbb{R}^{d+1}$ (resp. $\mathbb{L}^d$). We define their corresponding Lorentz Gramian (resp. H -Gramian) as

$$\begin{aligned} G &= ([x_i, x_j])_{i,j \in [N]} \\ &= X^T H X \end{aligned}$$

where H is the indefinite matrix given by (2).

The subtle difference between the Lorentz Gramian (defined for points in $\mathbb{R}^{d+1}$) and the H -Gramian (defined only on $\mathbb{L}^d \subset \mathbb{R}^{d+1}$) will be important for the low-rank projection and the spectral factorization steps in Section 3. The indefinite inner product (1) also determines the distance between $x, y \in \mathbb{L}^d$, as

$$d(x, y) = \text{acosh}(-[x, y]). \quad (4)$$

2.1.2 Poincaré Model. In the Poincaré model ($\mathbb{I}^d$), the points reside in the unit d -dimensional Euclidean ball. The distance between $x, y \in \mathbb{I}^d$ is given by

$$d(x, y) = \text{acosh} \left(1 + 2 \frac{\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)} \right). \quad (5)$$

The isometric map between the 'Loid and the Poincaré model, $h : \mathbb{L}^d \rightarrow \mathbb{I}^d$ is called the *stereographic projection*. For $y = h(x)$, we have

$$y_i = \frac{x_{i+1}}{x_0 + 1}, \quad (6)$$

The inverse of stereographic projection is given by

$$x = h^{-1}(y) = \frac{1}{1 - \|y\|^2} \begin{bmatrix} 1 + \|y\|^2 \\ 2y \end{bmatrix}. \quad (7)$$

The isometry between the 'Loid and Poincaré models makes them equivalent in their embedding capabilities. However, the Poincaré model facilitates visualization of the embedded points in a bounded disk, whereas the 'Loid model is an unbounded space.

2.2 Hyperbolic Distance Problems

In a metric hyperbolic distance problem, we want to find a point set $x_1, \dots, x_N \in \mathbb{L}^d$, such that

$$d_{mn} = \text{acosh}(-[x_m, x_n]), \quad \text{for all } (m, n) \in C,$$

for a subset of measured distances $\mathcal{D} = \{d_{mn} : (m, n) \in C \subseteq [N]_{\text{as}}^2\}$.

In many applications we have access to the *true* distances only through an unknown non-linear map $\tilde{d}_{mn} = \phi(d_{mn})$; examples are connectivity strength of neurons [22] or odor co-occurrence statistics [61]. If all we know is that $\phi(\cdot)$ is a monotonically increasing function, then only the ordinal information has remained intact,

$$d_{kl} \leq d_{mn} \Leftrightarrow \tilde{d}_{kl} \leq \tilde{d}_{mn}.$$

This leads to non-metric problems in which the measurements are in the form of binary comparisons [1].

DEFINITION 4. For a set of binary distance comparisons of the form $d_{kl} \leq d_{mn}$, we define the set of ordinal distance measurements as

$$\mathcal{O} = \{(k, l, m, n) : d_{kl} \leq d_{mn}, (k, l), (m, n) \in [N]_{\text{as}}^2\}.$$

We are now in a position to give a unified definition of metric and non-metric embedding problems in a hyperbolic space.

PROBLEM 1. A hyperbolic distance geometry problem aims to find $x_1, \dots, x_N \in \mathbb{L}^d$, given

- a subset of pairwise distances $\mathcal{D}$ such that

$$d_{mn} = d(x_m, x_n), \quad \text{for all } d_{mn} \in \mathcal{D}$$

- and/or a subset of ordinal distances measurements O such that

$$d(x_{i_1}, x_{i_2}) \leq d(x_{i_3}, x_{i_4}), \text{ for all } i \in O.$$

where $d(x, y) = \text{acosh}(-[x, y])$ and $i = (i_1, i_2, i_3, i_4)$.

We denote the complete sets of metric and non-metric measurements by $\mathcal{D}_c$ and O_c .

3 HYPERBOLIC DISTANCE MATRICES

We now introduce hyperbolic distance matrices in analogy with Euclidean distance matrices to compactly encode inter-point distances of a set of points $x_1, \dots, x_N \in \mathbb{L}^d$.

DEFINITION 5. The hyperbolic distance matrix (HDM) corresponding to the list of points $X = [x_1, \dots, x_N] \in (\mathbb{L}^d)^N$ is defined as

$$D = \mathcal{D}(X) = \left(d(x_i, x_j) \right)_{i,j \in [N]}.$$

The ij -th element of $\mathcal{D}(X)$ is hyperbolic distance between x_i and x_j , given by $d(x_i, x_j) = \text{acosh}(-[x_i, x_j])$ and for all $i, j \in [N]$.

HDMs are characterized by Lorentzian inner products which allows us to leverage the definition of an H-Gramian (Definition 3). Given points $x_1, \dots, x_N \in \mathbb{L}^d$, we compactly write the HDM corresponding to G as

$$D = \text{acosh}[-G], \quad (8)$$

where $\text{acosh}[\cdot]$ is an elementwise $\text{acosh}(\cdot)$ operator.

We now state our first main result: a semidefinite characterization of H -Gramians. This is a key step in casting HDGPs as rank-constrained semidefinite programs.

PROPOSITION 1 (SEMIDEFINITE CHARACTERIZATION OF H-GRAMIAN). Let G be the hyperbolic Gram matrix for a set of points $x_1, \dots, x_N \in \mathbb{L}^d$. Then,

$$G = G^+ - G^-$$

$$\begin{aligned} \text{where} \quad & G^+, G^- \geq 0, \\ & \text{rank } G^+ \leq d, \\ & \text{rank } G^- \leq 1, \\ & \text{diag } G = -1, \\ & e_i^\top G e_j \leq -1, \forall i, j \in [N]. \end{aligned}$$

Conversely, any matrix $G \in \mathbb{R}^{N \times N}$ that satisfies the above conditions is a hyperbolic Gramian for a set of N points in $\mathbb{L}^d$.

The proof is given in Proposition 1.

3.1 Solving for the H-Gramians

While Problem 1 could be formalized directly in X domain, this approach is unfavorable as the optimization domain, $\mathbb{L}^d$, is a non-convex set. What is more, the hyperbolic distances

$$d(x_m, x_n) = \text{acosh}(-e_m^\top X^\top H X e_n) \quad (9)$$

are non-linear functions of X with an unbounded gradient [13]. Similar issues arise when computing embeddings in other spaces such as Euclidean [16] or the space of polynomial trajectories [48]. A particularly effective strategy in the Euclidean case is the semidefinite relaxation which relies on the simple fact that the Euclidean

Gramian is positive semidefinite. We thus proceed by formulating a semidefinite relaxation for hyperbolic embeddings based on Proposition 1.

Solving the HDGP involves two steps, summarized in Algorithm 1:

- (1) Complete and denoise the HDM via a semidefinite program;
- (2) Compute an embedding of the clean HDM: we propose a closed-form spectral factorization method.

Note that step (2) is independent of step (1): given accurate hyperbolic distances, spectral factorization will give the points that reproduce them. However, since the semidefinite relaxation might give a Gramian with a higher rank than desired, eigenvalue thresholding in step (2) might move the points off of $\mathbb{L}^d$. That is because eigenvalue thresholding can violate the necessary condition for the hyperbolic norm, $\|x\|_H^2 = -1$, or $\text{diag } G = -1$ in Proposition 1. We fix this by projecting each individual point to $\mathbb{L}^d$. The spectral factorization and the projection are summarized in Algorithms 2 and 3.

Let $\tilde{D}$ be the measured noisy and incomplete HDM, with unknown entries replaced by zeroes. We define the mask matrix $W = (w_{ij})$ as

$$w_{ij} \stackrel{\text{def}}{=} \begin{cases} 1, & \text{for } (i, j) \in C \vee (j, i) \in C \\ 0, & \text{otherwise.} \end{cases}$$

This mask matrix lets us compute the loss only at those entries that were actually measured. We use the semidefinite characterization of hyperbolic Gramians in Proposition 1 to complete and denoise the measured HDM, and eventually solve HDGP.

Although the set of hyperbolic Gramians for a given embedding dimension is non-convex due to the rank constraints, discarding the rank constraints results in a straightforward semidefinite relaxation. However, if we convexify the problem by simply discarding the rank constraints, then all pairs $(G_1, G_2) \in \{(G^+ + P, G^- + P) : P \geq 0\}$ become a valid solution. On the other hand, since

$$\text{rank } G + P \geq \text{rank } G \text{ for } G, P \geq 0,$$

we can eliminate this ambiguity by promoting low-rank solutions for G^+ and G^- . While directly minimizing

$$\text{rank } G^+ + \text{rank } G^- \quad (10)$$

Algorithm 1 HDGP algorithm

- 1: **procedure** HDGP($\tilde{D}, \tilde{O}, d$)
 - 2: **Input:** Incomplete and noisy distance matrix, $\tilde{D}$, and ordinal measurements, $\tilde{O}$, and embedding dimension, d .
 - 3: $G = \text{SDR}(\tilde{D}, \tilde{O}, d)$ ▷ Complete & denoise HDM.
 - 4: $X = \text{Embed}(G, d)$ ▷ Embed points in $\mathbb{L}^d$.
 - 5: For $X = [x_1, \dots, x_N] \in (\mathbb{L}^d)^N$, let

$$y_n = h(x_n), \forall n \in [N]$$
 where $h(\cdot)$ is given by (6). ▷ Map the points to $\mathbb{I}^d$.
 - 6: **return** $Y = [y_1, \dots, y_N] \in (\mathbb{I}^d)^N$.
 - 7: **end procedure**
-

In Algorithm 3 we propose a simple but suboptimal procedure to solve this low-rank approximation problem. Unlike iterative refinement algorithms based on optimization on manifolds [25], our proposed method is one-shot. It is based on the spectral factorization of the estimated hyperbolic Gramian and involves the following steps:

- **Step 1:** We find a set of points $\{z_n\}$ in $\mathbb{R}^{d+1}$, whose Lorentz Gramian best approximates G ; See Definition 3 and lines 2 to 5 of Algorithm 3. In other words, we relax the optimization domain of (11) from $\mathbb{L}^d$ to $\mathbb{R}^{d+1}$,

$$Z = \arg \min_{X \in \mathbb{R}^{(d+1) \times N}} \|G - X^\top H X\|^2.$$

- **Step 2:** We project each point z_n onto $\mathbb{L}^d$, i.e.

$$\hat{X} = \arg \min_{X \in (\mathbb{L}^d)^N} \|X - Z\|_F^2.$$

This gives us an approximate rank- $(d+1)$ hyperbolic Gramian, $\hat{G} = \hat{X}^\top H \hat{X}$; see Figure 2 and Appendix C.

The first step of low-rank approximation of a hyperbolic Gramian G can be interpreted as finding the positions of points in $\mathbb{R}^{d+1}$ (not necessarily on $\mathbb{L}^d$) whose Lorentz Gramian best approximates G .

3.3 Spectral Factorization of H-Gramians

To finally compute the point locations, we describe a spectral factorization method, proposed in [55] (cf. footnote 1), to estimate point positions from their Lorentz Gramian (line 5 of Algorithm 3). This method exploits the fact that Lorentz Gramians have only one non-positive eigenvalue (see Lemma 1 in the appendix) as detailed in the following proposition.

PROPOSITION 2. *Let G be a hyperbolic Gramian for $X \in (\mathbb{L}^d)^N$, with eigenvalue decomposition $G = U\Lambda U^\top$, and eigenvalues $\lambda_0 \leq$*

Algorithm 3 Low-rank approximation and spectral factorization of hyperbolic Gramian

- 1: **procedure** Embed(G, d)
 - 2: **Input:** Hyperbolic Gramian G , and embedding dimension d .
 - 3: Let $U^\top \Lambda U$ be eigenvalue decomposition of G , where $\Lambda = \text{diag}(\lambda_0, \dots, \lambda_{N-1})$ such that
 - $\lambda_0 = \min_i \lambda_i$,
 - λ_i is the top i -th element of $\{\lambda_i\}$ for $i \in [N] - 1$.
 - 4: Let $G_{d+1} = U_d^\top \Lambda_d U_d$, where

$$\Lambda_d = \text{diag}(\lambda_0, u(\lambda_1), \dots, u(\lambda_d)),$$

$$u(x) = \max\{x, 0\}, \text{ and } U_d \text{ be the corresponding sliced eigenvalue matrix.}$$
 - 5: $Z = R|\Lambda_d|^{1/2}U_d^\top$, for any H-unitary matrix R .
 - 6: For $Z = [z_1, \dots, z_N]$, let

$$x_n = \text{Project}(z_n), \forall n \in [N]$$
 - 7: **return** $X = [x_1, \dots, x_N] \in (\mathbb{L}^d)^N$.
 - 8: **end procedure**
-

Figure 2: Projecting a point in $\mathbb{R}^{d+1}$ (blue) to $\mathbb{L}^d$ (red).

$0 \leq \lambda_1 \leq \dots \leq \lambda_d$.³ Then, there exists an H-unitary matrix R such that $X = R|\Lambda|^{1/2}U$.

The proof is given in Appendix D. Note that regardless of the choice of R , $X = R|\Lambda|^{1/2}U$ will reproduce G and thus the corresponding distances. This is the rigid motion ambiguity familiar from the Euclidean case [9]. If we start with an H-Gramian with a wrong rank, we need to follow the spectral factorization by Step 2 where we project each point $z_n \in \mathbb{R}^{d+1}$ onto $\mathbb{L}^d$. This heuristic is suboptimal, but it is nevertheless appealing since it only requires a single one-shot calculation as detailed in Appendix C.

4 EXPERIMENTAL RESULTS

In this section we numerically demonstrate different properties of Algorithm 1 in solving HDGPs. In a general hyperbolic embedding problem, we have a mix of metric and non-metric distance measurements which can be noisy and incomplete. Code, data and documentation to reproduce the experimental results are available at <https://github.com/puoya/hyperbolic-distance-matrices>.

4.1 Missing Measurements

Missing measurements are a common problem in hyperbolic embeddings of concept hierarchies. For example, hyperbolic embeddings of words based on Hearst-like patterns rely on co-occurrence probabilities of word pairs in a corpus such as WordNet [37]. These patterns are sparse since word pairs must be detected in the right configuration [31]. In perceptual embedding problems, we ask individuals to rate pairwise similarities for a set of objects. It may be difficult to collect and embed all pairwise comparisons in applications with large number of objects [1].

The proposed semidefinite relaxation gives a simple way to handle missing measurements. The *metric sampling density* $0 \leq S \leq 1$ of a measured HDM is the ratio of the number of missing measurements to total number of pairwise distances, $S = 1 - \frac{|\mathcal{D}|}{|\mathcal{D}_c|}$. We want to find the probability $p(S)$ of successful estimation given a sampling density S . In practice, we fix the embedding dimension, d , and the number of points, N , and randomly generate a point set, $X \in (\mathbb{L}^d)^N$. A trial is successful if we can solve the HDGP for noise-free measurements and a random mask W of a fixed size so that the estimated hyperbolic Gramian has a small relative error,

³An H-Gramian is a Lorentz Gramian.

Figure 3: Left and middle: The probability of δ -accurate estimation for metric sampling density S , $M = 100$, and $\delta = 10^{-2}$. Right: The empirical error $e_{\text{rel}} = \mathbb{E}_K[e_{\text{rel}}(X)]$ for ordinal sampling density S , $d = 2$, $M = 50$, and $K = 10$. In each bar, shading width represents the empirical standard deviation of $e_{\text{rel}}(X)$.

$e_{\text{rel}}(\hat{G}) = \frac{\|\mathcal{D}(X) - \text{acosh}[-\hat{G}]\|_F}{\|\mathcal{D}(X)\|_F} \leq \delta$. We repeat for M trials, and empirically estimate the success probability as $\hat{p}(S) = \frac{M_s}{M}$ where M_s is the number of successful trials. We repeat the experiment for different values of N and d , see Figure 3.

For non-metric embedding applications, we want to have *consistent* embedding for missing ordinal measurements. The *ordinal sampling density* $0 \leq S \leq 1$ of a randomly selected set of ordinal measurements is defined as $S = 1 - \frac{|O|}{|O_c|}$. For a point set $X \in (\mathbb{L}^d)^N$, we define the average relative error of estimated HDMs as $e_{\text{rel}}(X) = \mathbb{E}_M \frac{\|D_O - \mathbb{E}_M[D_O]\|_F}{\|\mathbb{E}_M[D_O]\|_F}$ where D_O is the estimated HDM for ordinal measurements O , and empirical expectation is with respect to the random ordinal set O . We repeat the experiment for K different realizations of $X \in (\mathbb{L}^d)^N$ (Figure 3). We can observe that across different embedding dimensions, the maximum allowed fraction of missing measurements for a consistent and accurate estimation increases with the number of points.

4.2 Weighted Tree Embedding

Tree-like hierarchical data occurs commonly in natural scenarios. In this section, we want to compare the embedding quality of weighted trees in hyperbolic and the baseline in Euclidean space.

We generate a random tree T with N nodes, maximum degree of $\Delta(T) = 3$, and i.i.d. edge weights from $\text{unif}(0, 1)^4$. Let D_T be the distance matrix for T , where the distance between each two nodes is defined as the weight of the path joining them.

For the hyperbolic embedding, we apply Algorithm 2 with log-det heuristic objective function to acquire a low-rank embedding. On the other hand, Euclidean embedding of T is the solution to the following semidefinite relaxation

$$\begin{aligned} & \text{minimize} && \|D_T^{\circ 2} - \mathcal{K}(G)\|_F^2 \\ & \text{w.r.t} && G \geq 0 \\ & \text{subject to} && G1 = 0 \end{aligned} \quad (12)$$

where $\mathcal{K}(G) = -2G + \text{diag}(G)1^\top + 1\text{diag}(G)^\top$, and $D_T^{\circ 2}$ is the entrywise square of D_T . This semidefinite relaxation (SDR) yields a *minimum error* embedding of T , since the embedded points can reside in an arbitrary dimensional Euclidean space.

The embedding methods based on semidefinite relaxation are generally accompanied by a projection step to account for the potentially incorrect embedding dimension. For hyperbolic embedding

problems, this step is summarized in Algorithm 3, whereas it is simply a singular value thresholding of the Gramian for Euclidean problems. Note that the SDRs always find a $(N - 1)$ -dimensional embedding for a set of N points; see Algorithm 2 and (12). In this experiment, we define the optimal embedding dimension as

$$d_0 = \min \left\{ d \in \mathbb{N} : \frac{\|D_{N-1} - D_d\|_F}{\|D_{N-1} - D_{d+1}\|_F} \geq 1 - \delta \right\}$$

where D_n is the distance matrix for embedded points in $\mathbb{L}^n$ (or $\mathbb{R}^n$), and $\delta = 10^{-3}$. This way, we accurately represent the estimated distance matrix in a low dimensional space. Finally, we define the relative (or normalized) error of embedding T in d_0 -dimensional

space as $e_{\text{rel}}(T) = \frac{\|D_T - D_{d_0}\|_F}{\|D_T\|_F}$. We repeat the experiment for M randomly generated trees T with a varying number of vertices N . The hyperbolic embedding yields smaller average relative error $\mathbb{E}_M[e_{\text{rel}}(T)]$ compared to Euclidean embedding, see Figure 4. It should also be noted that the hyperbolic embedding has a lower optimal embedding dimension, even though the low-rank hyperbolic Gramian approximation is sub-optimal.

4.3 Odor Embedding

In this section, we want to compare hyperbolic and Euclidean non-metric embeddings of olfactory data following the work of Zhou et al. [61]. We conduct identical experiments in each space, and compare embedding quality of points from Algorithm 2 in hyperbolic space to its semidefinite relaxation counterpart in Euclidean space, namely generalized non-metric MDS [1].

Figure 4: Tree embedding in hyperbolic (red) and Euclidean (green) space. Discrete distribution of optimal embedding dimension for $M = 100$, (a) and (b). Average, $\mathbb{E}_M[e_{\text{rel}}(T)]$, and standard deviation of embedding error, (c) and (d).

⁴The most likely maximum degree for trees with $N \leq 25$ [39].

Figure 5: Embedding of odors for different levels of allowable violated measurements ζ_p . Clusters with the matching colors contain the same odors.

We use an olfactory dataset comprising mono-molecular odor concentrations measured from blueberries [21]. In this dataset, there are $N = 52$ odors across the total of $M = 164$ fruit samples. Like Zhou *et al.* [61], we begin by computing correlations between odor concentrations across samples [61]. The correlation coefficient between two odors x_i and x_j is defined as

$$C(i, j) = \frac{(\mathbf{x}_i - \mu_{x_i} \mathbf{1})^\top (\mathbf{x}_j - \mu_{x_j} \mathbf{1})}{\|\mathbf{x}_i - \mu_{x_i} \mathbf{1}\| \|\mathbf{x}_j - \mu_{x_j} \mathbf{1}\|}$$

where $\mathbf{x}_n = (x_n^{(1)}, \dots, x_n^{(M)})^\top$, $x_i^{(m)}$ is the concentration of i -th odor in m -th fruit sample, M is total number of fruit samples and $\mu_{x_n} = \frac{1}{M} \sum_{m=1}^M x_n^{(m)}$.

The goal is to find an embedding for odors $y_1, \dots, y_N \in \mathbb{I}^d$ (or $\mathbb{R}^d$) such that

$$d(y_{i_1}, y_{i_2}) \leq d(y_{i_3}, y_{i_4}), (i_1, i_2, i_3, i_4) \in \mathcal{O},$$

where,

$$\mathcal{O} \subseteq \mathcal{O}_c = \left\{ (i_1, i_2, i_3, i_4) \in ([N]_{\text{as}}^2)^2 : C(i_1, i_2) \geq C(i_3, i_4) \right\}.$$

The total number of distinct comparisons grows rapidly with the number of points, namely $|\mathcal{O}_c| = 0.87$ million. In this experiment, we choose a random set of size $|\mathcal{O}| = 2K \binom{N}{2}$ for $K = 4$ to have the sampling density of $S = 98.79\%^5$, which brings the size of ordinal measurements to $|\mathcal{O}| \approx 10^4$.

We ensure the embedded points do not collapse by imposing the following minimum distance constraint $d(x_i, x_j) \geq 1$ for all $(i, j) \in [N]_{\text{as}}^2$; this corresponds to a simple linear constraint in the proposed formulation. An ideal order embedding accurately reconstructs the missing comparisons. We calculate the percentage of correctly reconstructed distance comparisons as $\gamma_d = |\widehat{\mathcal{O}}_{c,d} \cap \mathcal{O}_c| / |\mathcal{O}_c|$, where $\widehat{\mathcal{O}}_{c,d}$ is the complete ordinal set corresponding to a d -dimensional embedding.

A simple regularization technique helps to remove outlier measurements and improve the generalized accuracy of embedding algorithms. We introduce the parameter ζ_p to permit SDR algorithms to dismiss at most p -percent of measurements, namely

$$\mathcal{L}_k(G) + \epsilon_k \geq \epsilon_2 \text{ and } \epsilon_k \geq 0, \forall k \in \mathcal{O} \text{ and } \sum_k \epsilon_k \leq \zeta_p$$

⁵In hyperbolic embedding, this is the ratio of number of ordinal measurements to number of variables, i.e. $K = \frac{|\mathcal{O}|}{2 \binom{N}{2}}$.

Table 4: Reconstruction accuracy of ordinal measurements γ_d for different levels of allowable violation ζ_p .

Space		$d = 2$	$d = 4$	$d = 6$	$d = 8$	$d = 10$
Hyperbolic	ζ_0	76.06	83.60	86.87	89.48	91.03
	$\zeta_{0.5}$	76.52	83.71	86.94	89.68	91.16
	ζ_1	76.43	83.71	86.92	89.76	91.21
Euclidean	ζ_0	73.44	78.86	82.23	85.06	88.67
	$\zeta_{0.5}$	73.27	79.03	82.65	86.24	88.98
	ζ_1	73.12	78.92	82.51	86.01	89.02

where $\zeta_p = \frac{p}{100} |\mathcal{O}| \epsilon_2$.

In Figure 5, we show the embedded points in $\mathbb{I}^2$ and $\mathbb{R}^2$ with different levels of allowable violated measurements. We can observe in Table 4 that hyperbolic space better represent the structure of olfactory data compared to Euclidean space of the same dimension. This is despite the fact that the number of measurements per variable is in favor of Euclidean embedding, and that the low rank approximation of hyperbolic Gramians is suboptimal. Moreover, if we remove a small number of outliers we can produce more accurate embeddings. These results corroborate the statistical analysis of Zhou *et al.* [61] that aims to identify the geometry of the olfactory space.⁶

5 CONCLUSION

We introduced hyperbolic distance matrices, an analogy to Euclidean distance matrices, to encode pairwise distances in the 'Loid model of hyperbolic geometry. Same as in the Euclidean case, although the definition of hyperbolic distance matrices is trivial, analyzing their properties gives rise to powerful algorithms based on semidefinite programming. We proposed a semidefinite relaxation which is essentially plug-and-play: it easily handles a variety of metric and non-metric constraints, outlier removal, and missing information and can serve as a template for different applications. Finally, we proposed a closed-form spectral factorization algorithm to estimate the point position from hyperbolic Gramians. Several important questions are still left open, most notably the role of the isometries in the 'Loid model and the related concepts such as Procrustes analysis.

ACKNOWLEDGEMENT

We thank Lav Varshney for bringing our attention to hyperbolic geometry and for the numerous discussions about the manuscript.

REFERENCES

- [1] Sameer Agarwal, Josh Wills, Lawrence Cayton, Gert Lanckriet, David Kriegman, and Serge Belongie. 2007. Generalized non-metric multidimensional scaling. In *Artificial Intelligence and Statistics*. 11–18.
- [2] Leman Akoglu, Mary McGlohon, and Christos Faloutsos. 2010. Oddball: Spotting anomalies in weighted graphs. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 410–421.
- [3] Abdo Y Alfakih, Amir Khandani, and Henry Wolkowicz. 1999. Solving Euclidean distance matrix completion problems via semidefinite programming. *Computational Optimization and Applications* 12, 1-3 (1999), 13–30.
- [4] Michael Ashburner, Catherine A Ball, Judith A Blake, David Botstein, Heather Butler, J Michael Cherry, Allan P Davis, Kara Dolinski, Selina S Dwight, Janan T

⁶Statistical analysis of Betti curve behavior of underlying clique topology [22].

- Eppig, et al. 2000. Gene ontology: tool for the unification of biology. *Nature Genetics* 25, 1 (2000), 25.
- [5] Dena Asta and Cosma Rohilla Shalizi. 2014. Geometric network comparison. *arXiv preprint arXiv:1411.1350* (2014).
- [6] Riccardo Benedetti and Carlo Petronio. 2012. *Lectures on hyperbolic geometry*. Springer Science & Business Media.
- [7] Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. 2010. Sustaining the Internet with hyperbolic mapping. *Nature Communications* 1 (2010), 62.
- [8] Carlo Vittorio Cannistraci, Gregorio Alanis-Lobato, and Timothy Ravasi. 2013. From link-prediction in brain connectomes and protein interactomes to the local-community-paradigm in complex networks. *Scientific Reports* 3 (2013), 1613.
- [9] James W Cannon, William J Floyd, Richard Kenyon, Walter R Parry, et al. 1997. Hyperbolic geometry. *Flavors of Geometry* 31 (1997), 59–115.
- [10] Benjamin Paul Chamberlain, Stephen R Hardwick, David R Wardrope, Fabon Dzugang, Fabio Daolio, and Saúl Vargas. 2019. Scalable hyperbolic recommender systems. *arXiv preprint arXiv:1902.08648* (2019).
- [11] Kenny Chowdhary and Tamara G Kolda. 2018. An improved hyperbolic embedding algorithm. *Journal of Complex Networks* 6, 3 (2018), 321–341.
- [12] Andrej Cvetkovski and Mark Crovella. 2009. Hyperbolic embedding and routing for dynamic graphs. In *IEEE International Conference on Computer Communications*. IEEE, 1647–1655.
- [13] Christopher De Sa, Albert Gu, Christopher Ré, and Frederic Sala. 2018. Representation tradeoffs for hyperbolic embeddings. *Proceedings of Machine Learning Research* 80 (2018), 4460.
- [14] Bhuwan Dhingra, Christopher J Shallue, Mohammad Norouzi, Andrew M Dai, and George E Dahl. 2018. Embedding text in hyperbolic spaces. *arXiv preprint arXiv:1806.04313* (2018).
- [15] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research* 17, 1 (2016), 2909–2913.
- [16] Ivan Dokmanić, Reza Parhizkar, Juri Ranieri, and Martin Vetterli. 2015. Euclidean distance matrices: Essential theory, algorithms, and applications. *IEEE Signal Processing Magazine* 32, 6 (2015), 12–30.
- [17] Maryam Fazel. 2002. Matrix rank minimization with applications. (2002).
- [18] Maryam Fazel, Haitham Hindi, and Stephen P Boyd. 2003. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003.*, Vol. 3. IEEE, 2156–2162.
- [19] Massimo Fornasier, Holger Rauhut, and Rachel Ward. 2011. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM Journal on Optimization* 21, 4 (2011), 1614–1640.
- [20] Octavian-Eugen Ganeva, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic entailment cones for learning hierarchical embeddings. *arXiv preprint arXiv:1804.01882* (2018).
- [21] Jessica L Gilbert, Matthew J Guthart, Salvador A Gezan, Melissa Pizaroglo de Carvalho, Michael L Schwieterman, Thomas A Colquhoun, Linda M Bartoshuk, Charles A Sims, David G Clark, and James W Olmstead. 2015. Identifying breeding priorities for blueberry flavor using biochemical, sensory, and genotype by environment analyses. *PLoS One* 10, 9 (2015), e0138494.
- [22] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. 2015. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences* 112, 44 (2015), 13455–13460.
- [23] Israel Gohberg, Peter Lancaster, and Leiba Rodman. 1983. Matrices and indefinite scalar products. (1983).
- [24] Roger A Horn and Charles R Johnson. 2012. *Matrix analysis*. Cambridge University Press.
- [25] Pratik Jawanpuria, Mayank Meghwanshi, and Bamdev Mishra. 2019. Low-rank approximations of hyperbolic embeddings. *arXiv preprint arXiv:1903.07307* (2019).
- [26] Martin Keller-Ressel and Stephanie Nargang. 2020. Hydra: a method for strain-minimizing hyperbolic embedding of network-and distance-based data. *Journal of Complex Networks* 8, 1 (2020), cnaa002.
- [27] Robert Kleinberg. 2007. Geographic routing using hyperbolic space. In *26th IEEE International Conference on Computer Communications*. IEEE, 1902–1909.
- [28] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106.
- [29] Joseph B Kruskal and Myron Wish. 1978. *Multidimensional scaling*. Number 11. Sage.
- [30] John Lamping and Ramana Rao. 1994. Laying out and visualizing large trees using a hyperbolic space. In *Proceedings of the 7th annual ACM symposium on User interface software and technology*. ACM, 13–14.
- [31] Matt Le, Stephen Roller, Laetitia Papaxanthos, Douwe Kiela, and Maximilian Nickel. 2019. Inferring concept hierarchies from text corpora via hyperbolic embeddings. *arXiv preprint arXiv:1902.00913* (2019).
- [32] Leo Liberti, Carlile Lavor, Nelson Maculan, and Antonio Mucherino. 2014. Euclidean distance geometry and applications. *SIAM Rev.* 56, 1 (2014), 3–69.
- [33] Nathan Linial, Eran London, and Yuri Rabinovich. 1995. The geometry of graphs and some of its algorithmic applications. *Combinatorica* 15, 2 (1995), 215–245.
- [34] Ke Ma, Qianqian Xu, and Xiaochun Cao. 2019. Robust ordinal embedding from contaminated relative comparisons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7908–7915.
- [35] Anirudha Majumdar, Georgina Hall, and Amir Ali Ahmadi. 2019. Recent scalability improvements for semidefinite programming with applications in machine Learning, control, and robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 3 (2019).
- [36] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.
- [37] George A Miller. 1998. *WordNet: An electronic lexical database*. MIT press.
- [38] Bamdev Mishra, Gilles Meyer, Francis Bach, and Rodolphe Sepulchre. 2013. Low-rank optimization with trace norm penalty. *SIAM Journal on Optimization* 23, 4 (2013), 2124–2149.
- [39] John W Moon et al. 1968. On the maximum degree in a random tree. *The Michigan Mathematical Journal* 15, 4 (1968), 429–432.
- [40] Maximilian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*. 6338–6347.
- [41] Maximilian Nickel and Douwe Kiela. 2018. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. *arXiv preprint arXiv:1806.03417* (2018).
- [42] Carl Olsson, Anders Eriksson, and Richard Hartley. 2010. Outlier removal using duality. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 1450–1457.
- [43] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.
- [44] Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. Hearst patterns revisited: Automatic hypernym detection from large text corpora. *arXiv preprint arXiv:1806.03191* (2018).
- [45] Rik Sarkar. 2011. Low distortion delaunay embedding of trees in hyperbolic plane. In *International Symposium on Graph Drawing*. Springer, 355–366.
- [46] Yongduek Seo, Hyunjung Lee, and Sang Wook Lee. 2009. Outlier removal by convex optimization for l-infinity approaches. In *Pacific-Rim Symposium on Image and Video Technology*. Springer, 203–214.
- [47] Yuval Shavit and Tomer Tankel. 2008. Hyperbolic embedding of internet graph for distance estimation and overlay construction. *IEEE/ACM Transactions on Networking* 16, 1 (2008), 25–36.
- [48] Puoya Tabaghi, Ivan Dokmanić, and Martin Vetterli. 2019. Kinetic Euclidean distance matrices. *IEEE Transactions on Signal Processing* 68 (2019), 452–465.
- [49] Omer Tamuz, Ce Liu, Serge Belongie, Ohad Shamir, and Adam Tauman Kalai. 2011. Adaptively learning the crowd kernel. *arXiv preprint arXiv:1105.1033* (2011).
- [50] Laurens Van Der Maaten and Kilian Weinberger. 2012. Stochastic triplet embedding. In *2012 IEEE International Workshop on Machine Learning for Signal Processing*. IEEE, 1–6.
- [51] Lieven Vandenbergh and Stephen Boyd. 1996. Semidefinite programming. *SIAM Rev.* 38, 1 (1996), 49–95.
- [52] Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2015. Order-embeddings of images and language. *arXiv preprint arXiv:1511.06361* (2015).
- [53] Kevin Verbeek and Subhash Suri. 2014. Metric embedding, hyperbolic space, and social networks. In *Proceedings of the thirtieth annual symposium on Computational Geometry*. ACM, 501.
- [54] Tran Dang Quang Vinh, Yi Tay, Shuai Zhang, Gao Cong, and Xiao-Li Li. 2018. Hyperbolic recommender systems. *arXiv preprint arXiv:1809.01703* (2018).
- [55] Richard C Wilson, Edwin R Hancock, Elzbieta Pekalska, and Robert PW Duin. 2014. Spherical and hyperbolic embeddings of data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 11 (2014), 2255–2269.
- [56] Huan Xu, Constantine Caramanis, and Sujay Sanghavi. 2010. Robust PCA via outlier pursuit. In *Advances in Neural Information Processing Systems*. 2496–2504.
- [57] Jin Yu, Anders Eriksson, Tat-Jun Chin, and David Suter. 2014. An adversarial optimization approach to efficient outlier removal. *Journal of Mathematical Imaging and Vision* 48, 3 (2014), 451–466.
- [58] Wenchao Yu, Wei Cheng, Charu C Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. 2018. Network: A flexible deep embedding approach for anomaly detection in dynamic networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2672–2681.
- [59] Alp Yurtsever, Joel A Tropp, Olivier Fercoq, Madeleine Udell, and Volkan Cevher. 2019. Scalable semidefinite programming. *arXiv preprint arXiv:1912.02949* (2019).
- [60] Alp Yurtsever, Madeleine Udell, Joel A Tropp, and Volkan Cevher. 2017. Sketchy decisions: Convex low-rank matrix optimization with optimal storage. *arXiv preprint arXiv:1702.06838* (2017).
- [61] Yuansheng Zhou, Brian H Smith, and Tatyana O Sharpee. 2018. Hyperbolic geometry of the olfactory space. *Science Advances* 4, 8 (2018), eaq1458.

A PROOF OF PROPOSITION 1

A hyperbolic Gramian can be written as $G = X^T H X$ for a $X = [x_1, \dots, x_N] \in (\mathbb{L}^d)^N$. Let us rewrite it as

$$\begin{aligned} G &= \sum_{i=1}^d g_i g_i^T - g_0 g_0^T \\ &= G^+ - G^- \end{aligned}$$

where g_i^T is the $(i+1)$ -th row of X , $G^- = g_0 g_0^T$ and $G^+ = \sum_{i=1}^d g_i g_i^T$ are positive semidefinite matrices. We have $\text{rank } G^- \leq 1$ and $\text{rank } G^+ \leq d$. On the other hand, we have

$$\begin{aligned} e_i^T G e_j &\stackrel{\text{def}}{=} [x_i, x_j] \\ &= -x_{0,i} x_{0,j} + \sum_{k=1}^d x_{k,i} x_{k,j} \\ &\stackrel{(a)}{=} -\sqrt{1 + \|\bar{x}_i\|^2} \sqrt{1 + \|\bar{x}_j\|^2} + \bar{x}_i^T \bar{x}_j \\ &\stackrel{(b)}{\leq} -(1 + \bar{x}_i^T \bar{x}_j) + \bar{x}_i^T \bar{x}_j = -1. \end{aligned}$$

where $x_{k,i}$ is the $(k+1)$ -th element of x_i , $\bar{x}_i = (x_{1,i}, \dots, x_{d,i})^T$, and (a) is due to $\|x_i\|_H^2 = \|x_j\|_H^2 = -1$, and (b) results from Cauchy-Schwartz inequality. The equality holds for $i = j$, which yields the $\text{diag } G = -1$ condition.

Conversely, let $G = G^+ - G^-$, where $G^+, G^- \geq 0$, $\text{rank } G^- \leq 1$, and $\text{rank } G^+ \leq d$. Let us write $G^- = g_0 g_0^T$ and $G^+ = \sum_{i=1}^d g_i g_i^T$ for $g_0, \dots, g_d \in \mathbb{R}^N$. Then, we define

$$X \stackrel{\text{def}}{=} \begin{bmatrix} g_0^T \\ \vdots \\ g_d^T \end{bmatrix} = [x_1, \dots, x_N] \in \mathbb{R}^{(d+1) \times N}.$$

where $x_n \in \mathbb{R}^{d+1}$ for all $n \in [N]$. By construction, we have $X^T H X = G$, and

$$\text{diag } G = -1 \Rightarrow \|x_n\|_H^2 = -1, \forall n \in [N].$$

Finally, $e_i^T G e_j \leq -1$ guarantees that $x_n \in \mathbb{L}^d$ for all $n \in [N]$. We prove the contrapositive statement. Let x_i and x_j belong to different the hyperbolic sheets, e.g. $x_i \in \mathbb{L}^d, x_j \in -\mathbb{L}^d$. Then,

$$\begin{aligned} e_i^T G e_j &\stackrel{\text{def}}{=} [x_i, x_j] \\ &= -x_{0,i} x_{0,j} + \sum_{k=1}^d x_{k,i} x_{k,j} \\ &\stackrel{(a)}{\geq} \sqrt{1 + \|\bar{x}_i\|^2} \sqrt{1 + \|\bar{x}_j\|^2} - \|\bar{x}_i\| \|\bar{x}_j\| \geq 0 \end{aligned}$$

where (a) is due to Cauchy-Schwartz inequality. This is in contradiction with $e_i^T G e_j \leq -1$ condition. Therefore, $\{x_n\}$ belong to the same hyperbolic sheet, namely $\mathbb{L}^d$.

B DERIVATIONS FOR ALGORITHM 3

THEOREM 1. Let $G \in \mathbb{R}^{N \times N}$ be a hyperbolic Gramian, with eigenvalue decomposition

$$G = U^T \Lambda U, \quad (13)$$

where $\Lambda = \text{diag}(\lambda_0, \dots, \lambda_{N-1})$ such that

- $\lambda_0 = \min_i \lambda_i$,
- λ_i is the i -th top element of $\{\lambda_i\}$ for $i \in \{1, \dots, d\}$

The best rank- $(d+1)$ Lorentz Gramian approximation of G , in ℓ_2 sense, is given by

$$G_{d+1} = U_d^T \Lambda_d U_d$$

where $\Lambda_d = \text{diag}[\lambda_0, u(\lambda_1), \dots, u(\lambda_d)]$, $u(x) = \max\{x, 0\}$, and $U_d \in \mathbb{R}^{(d+1) \times N}$ is the corresponding sliced eigenvalue matrix.

PROOF. We begin by characterizing the eigenvalues of a Lorentz Gramian.

LEMMA 1. Let $G \in \mathbb{R}^{N \times N}$ be a Lorentz Gramian of rank $d+1$ with eigenvalues $\psi_0 \leq \dots \leq \psi_d$. Then, $\psi_0 < 0$, and $\psi_i > 0$, for $i \in \{1, \dots, d\}$.

PROOF. We write Lorentzian Gramian, $G = (g_{i,j})$, as $G = X^T H X$ where

$$X = [x_1, \dots, x_N] \stackrel{\text{def}}{=} \begin{bmatrix} g_0^T \\ \vdots \\ g_d^T \end{bmatrix} \in \mathbb{R}^{(d+1) \times N}.$$

Then, $G = G^+ - G^-$ where $G^+ \stackrel{\text{def}}{=} \sum_{i=1}^d g_i g_i^T$ is a positive semidefinite matrix of rank d and with eigenvalues $0 < \gamma_1 \leq \dots \leq \gamma_d$, and $-G^- \stackrel{\text{def}}{=} -g_0 g_0^T$ is a negative definite matrix of rank 1, with eigenvalue $\mu \leq 0$. From Weyl's inequality [24], we have

$$\mu + \gamma_1 \leq \psi_0 \leq \mu + \gamma_d$$

where ψ_0 is the smallest eigenvalue of G . Therefore, ψ_0 can be non-positive (negative if $\mu + \gamma_d < 0$). For other eigenvalues of G , we have

$$0 + \gamma_1 \leq \psi_i \leq \gamma_d, \text{ for } 1 \leq i \leq d.$$

Hence, $\psi_i > 0$ for $i \in \{1, \dots, d\}$. This result is irrespective to the order of eigenvalues.

Now, let us prove $\psi_0 < 0$. Suppose $g_0 \in S = \text{span}\{g_i : i \in \{1, \dots, d\}\}$. Then,

$$\text{rank } G = \text{rank} \begin{bmatrix} g_0^T \\ \vdots \\ g_d^T \end{bmatrix} < d+1,$$

which is a contradiction. Therefore, we write $g_0 = \alpha t + \beta s$ where $s \in S$, $t \in S^\perp$ with $\|t\| = 1$, $\alpha, \beta \in \mathbb{R}$ and $\alpha \neq 0$. Then, we have

$$\begin{aligned} \psi_0 &\leq t^T G t \\ &\stackrel{(a)}{=} -t^T g_0 g_0^T t \\ &= -\alpha^2 < 0 \end{aligned}$$

where (a) is due to $G = -g_0 g_0^T + \sum_{i=1}^d g_i g_i^T$ and $t \in S^\perp$. $\square$

Consider eigenvalue decomposition of G in eq. (13). Without loss of generality, we assume

- $\lambda_0 = \min_i \lambda_i < 0$,
- λ_i is the i -th top element of $\{\lambda_i\}$ for $i \in \{1, \dots, d\}$.

By construction $G = X^T H X$ and from $\text{diag } G = -1$ condition, we have

$$\sum \lambda_i = -N.$$

Therefore, $\lambda_0 < 0$. From Lemma 1, one eigenvalue of a Lorentz Gramian is negative and the rest must be positive. Therefore, $\hat{G} =$

$U_d^\top \Lambda_d U_d$ with eigenvalues $\Lambda_d = \text{diag}\{\lambda_0, u(\lambda_1), \dots, u(\lambda_d)\}$ and eigenvectors $U_d = [u_0, \dots, u_d]$, is the best rank- $(d+1)$ Lorentz Gramian approximation to G , i.e.

$$\|\hat{G} - G\|_2^2 = \inf_{H: \text{Lorentz Gram. of rank } \leq d+1} \|H - G\|_2^2.$$

□

Finally, a rank- $(d+1)$ Lorentz Gramian with eigenvalue decomposition

$$G_{d+1} = U_d \Lambda_d U_d^\top$$

can be decomposed as $X = R|\Lambda|^{1/2}U_d^\top \in \mathbb{R}^{(d+1) \times N}$ where R is an arbitrary H-unitary matrix and $G_{d+1} = X^\top H X$.

C Project : $\mathbb{R}^d \rightarrow \mathbb{L}^d$

Algorithm 4 Projection from $\mathbb{R}^{d+1}$ to $\mathbb{L}^d$

1: **procedure** Project(x)

2: For $x \in \mathbb{R}^{d+1}$, let

$$\hat{x} = \begin{cases} (1, 0^\top)^\top & x \in \{(x_0, 0^\top)^\top : x_0 \leq 2\}, \\ (\frac{1}{2}x_0, \hat{x}_1, \dots, \hat{x}_d)^\top & x \in \{(x_0, 0^\top)^\top : x_0 > 2\} \\ & \text{and for a } (\hat{x}_1, \dots, \hat{x}_d) \in S, \\ x(\lambda) & \text{otherwise and for } \lambda : \\ & \|x(\lambda)\|_H^2 = -1. \end{cases}$$

where $x(\lambda) = (I + \lambda H)^{-1}x$ and

$$S = \left\{ (x_1, \dots, x_d) : x_1^2 + \dots + x_d^2 = -1 + \frac{1}{4}x_0^2 \right\}.$$

3: **return** $\hat{x}$.

4: **end procedure**

PROOF. Let us reformulate the following projection problem

$$\hat{x} \in \arg \min_{y \in \mathbb{L}^d} \|y - x\|^2 \quad (14)$$

as unconstrained augmented Lagrangian minimization, i.e.

$$L(y, \lambda) = \|y - x\|^2 + \lambda(y^\top H y + 1).$$

The first order necessary condition for $\hat{x}$ to be a (local) minimum of eq. (14) is

$$(I + \lambda^* H)\hat{x} = x \quad (15)$$

for a $\lambda^* \in \mathbb{R}$ such that $\hat{x} \in \mathbb{L}^d$.

$\lambda^* = -1$: This happens when $x = (x_0, 0^\top)^\top$ and $x_0 \geq 2$. Following from optimality condition of eq. (15) and $\|\hat{x}\|_H^2 = -1$, we have $\hat{x} = (\frac{1}{2}x_0, \hat{x}_1, \dots, \hat{x}_d)^\top$, where

$$\hat{x}_1^2 + \dots + \hat{x}_d^2 = -1 + \frac{1}{4}x_0^2.$$

Therefore, $\hat{x}$ could be any point on a $(d-1)$ -dimensional sphere on $\mathbb{L}^d$. For $x = (x_0, 0^\top)^\top$ and $x_0 \leq 2$, we have $\hat{x} = (1, 0^\top)^\top$.

$\lambda^* = 1$: This happens for $x = (0, x_1, \dots, x_d)^\top$. From optimality condition of eq. (15), we have $\hat{x} = (\hat{x}_0, \frac{1}{2}x_1, \dots, \frac{1}{2}x_d)$, where $\hat{x}_0 = \frac{1}{2}\sqrt{x_1^2 + \dots + x_d^2 + 4}$.

For non-degenerate cases of $\lambda^* \neq \pm 1$, we have

$$\hat{x} = (I + \lambda^* H)^{-1}x, \quad (16)$$

where $\lambda^* \in \left\{ \lambda : \|(I + \lambda H)^{-1}x\|_H^2 = -1, \hat{x}_0 \geq 0 \right\}$.

(1) $\lambda^* \in (-1, 1)$: First, we define

$$f(\lambda) = \|(I + \lambda H)^{-1}x\|_H^2.$$

This is a monotonous function on $(-1, 1)$, with $\lim_{\lambda \rightarrow -1^-} f(\lambda) = -\infty$, and $\lim_{\lambda \rightarrow -1^+} f(\lambda) = +\infty$. Hence, $f(\lambda) = -1$ has a unique solution $\lambda^* \in (-1, 1)$. Finally, $\hat{x}$ is a local minima since the second order sufficient condition

$$I + \lambda^* H > 0$$

is satisfied for $\lambda^* \in (-1, 1)$. Lastly, from eq. (16), we have $\hat{x}_0 x_0 \geq 0$. In other words, $\lambda^* \in [-1, 1]$ if and only if x is in the same half-space as $\mathbb{L}^d$, i.e. $x_0 \geq 0$.

(2) $\lambda^* \in (-\infty, -1)$: Similarly, $f(\lambda)$ is a continuous function in this interval with $\lim_{\lambda \rightarrow -1^-} f(\lambda) = +\infty$, $\lim_{\lambda \rightarrow -\infty} f(\lambda) = 0$, and its first order derivative

$$\frac{d}{d\lambda} f(\lambda) = -\frac{2}{(1-\lambda)^3} x_0^2 - \frac{2}{(1+\lambda)^3} \sum_{i=1}^d x_i^2$$

has at most one zero. Therefore, $f(\lambda) = -1$ has at most two solutions. The second order necessary condition for local minima is $v^\top (I + \lambda^* H)v \geq 0$ for all $v \in T_{\hat{x}} \mathbb{L}^d$, where

$$T_{\hat{x}} \mathbb{L}^d = \{v \in \mathbb{R}^{d+1} : x^\top (I + \lambda^* H)^{-1} H v = 0\}.$$

However, there exists a $v \in T_{\hat{x}} \mathbb{L}^d$ where $v = (0, \bar{v}^\top)^\top$ which violates the second order necessary condition, $v^\top (I + \lambda^* H)v < 0$. Therefore, $\hat{x}$ – even if it exists – is not a local minima.

(3) $\lambda^* \in (1, \infty)$: We can easily see that $\lim_{\lambda \rightarrow 1^+} f(\lambda) = -\infty$, $\lim_{\lambda \rightarrow +\infty} f(\lambda) = 0$, and $\frac{d}{d\lambda} f(\lambda) = 0$ has at most one solution in this interval. Therefore, $f(\lambda) = -1$ has exactly one solution. However, we have $\hat{x}_0 x_0 \leq 0$ from eq. (16). In other words, $\lambda^* \in (1, \infty)$ if and only if x is in the opposite half-space of $\mathbb{L}^d$, i.e. $x_0 \leq 0$. Finally, $\hat{x}$ is the unique minima, since the projection of $x \notin S$ to the closed and convex set of

$$S = \{x : x_0 \geq 0, \|x\|_H^2 \leq -1\}$$

always exists and is unique. □

D PROOF OUTLINE OF PROPOSITION 2

Let $X = R|\Lambda|^{1/2}U^\top$. Then,

$$X^\top H X = U|\Lambda|^{1/2}R^\top H R|\Lambda|^{1/2}U^\top$$

$$\stackrel{(a)}{=} U|\Lambda|^{1/2}H|\Lambda|^{1/2}U^\top$$

$$\stackrel{(b)}{=} G$$

where (a) is due to properties of H-unitary matrices, (b) from $|\lambda_0|^{1/2}(-1)|\lambda_0|^{1/2} = \lambda_0$ for $\lambda_0 \leq 0$. Therefore $X = R|\Lambda|^{1/2}U^\top$ is a hyperbolic spectral factor of G . Finally, the uniqueness of these factors is due to fact that H -unitary operators fully characterize isometric maps in the 'Loid model [9].