
Hyperbolic Procrustes Analysis Using Riemannian Geometry

Ya-Wei Eileen Lin[†] Yuval Kluger^{‡§¶} Ronen Talmon[†]

[†]Viterbi Faculty of Electrical and Computer Engineering, Technion

[‡]Program in Applied Mathematics, Yale University

[§]Interdepartmental Program in Computational Biology and Bioinformatics, Yale University

[¶]Department of Pathology, Yale University

{lin.ya-wei@campus, ronent@ee}.technion.ac.il, yuval.kluger@yale.edu

Abstract

Label-free alignment between datasets collected at different times, locations, or by different instruments is a fundamental scientific task. Hyperbolic spaces have recently provided a fruitful foundation for the development of informative representations of hierarchical data. Here, we take a purely geometric approach for label-free alignment of hierarchical datasets and introduce hyperbolic Procrustes analysis (HPA). HPA consists of new implementations of the three prototypical Procrustes analysis components: translation, scaling, and rotation, based on the Riemannian geometry of the Lorentz model of hyperbolic space. We analyze the proposed components, highlighting their useful properties for alignment. The efficacy of HPA, its theoretical properties, stability and computational efficiency are demonstrated in simulations. In addition, we showcase its performance on three batch correction tasks involving gene expression and mass cytometry data. Specifically, we demonstrate high-quality unsupervised batch effect removal from data acquired at different sites and with different technologies that outperforms recent methods for label-free alignment in hyperbolic spaces.

1 Introduction

A key scientific task in modern data analysis is the alignment of data. The need for alignment often arises since data are acquired in multiple domains, under different environmental conditions, using various acquisition equipment, and at different sites. This paper focuses on the problem of label-free alignment of data embedded in hyperbolic spaces. Recently, hyperbolic spaces have accentuated in geometric representation learning. These non-Euclidean spaces have become popular since they provide a natural embedding of hierarchical data thanks to the exponential growth of the lengths of their geodesic paths [41, 42, 30, 15, 14, 6, 32].

The problem of alignment of data embedded in hyperbolic spaces has been extensively studied, e.g., in the context of natural language processing [49], ontology matching [10], matching two data modalities [40], and improving the embedding in hyperbolic spaces [2]. A few of these studies are based on optimal transport (OT) [2, 22], a classical problem in mathematics [38] that has recently reemerged in modern data analysis, e.g., for domain adaptation [7]. Despite its increasing usage, OT for unsupervised alignment is fundamentally limited [54], since OT (as any density matching approach) cannot recover volume-preserving maps [3, 4, 36].

In this paper, we resort to Procrustes analysis (PA) [17, 18] that is based on purely geometric considerations. PA has been widely used for aligning datasets by eliminating the shift, scaling, and rotational factors. Over the years, it has been successfully applied to various applications, e.g.,

image registration [34], manifold alignment [52], shape matching [35], domain adaptation [47], and manifold learning [27], to name but a few. Here, we address the problem of label-free matching of hierarchical data embedded in hyperbolic spaces. We present hyperbolic Procrustes analysis (HPA), a new PA method in the Lorentz model of hyperbolic geometry. The main novelty lies in the introduction of new implementations of the three prototypical PA components based on Riemannian geometry. Specifically, translation is viewed as a Riemannian mean alignment, implemented using parallel transport (PT). Scaling is determined with respect to geodesic paths. Rotation is considered as moment alignment on a mapping of the tangent space of the manifold to a Euclidean vector space. Our analysis provides new derivations in the Riemannian geometry of the Lorentz model and specifies the commuting properties of the HPA components. We show that HPA, compared to existing baselines and OT-based methods, achieves improved alignment in a purely unsupervised setting. In addition, it has a natural and stable out-of-sample extension, it supports both small and big data, and it is computationally efficient.

We show application to batch correction in bioinformatics tasks. We present results on both gene expression and mass cytometry (CyTOF) data, exemplifying the generality and broad scope of our method. In contrast to recent works [28, 50], our method does not require landmark correspondence, which is often unavailable in many datasets or hard to obtain. Specifically, we show that batch effects caused by acquisition using different technologies, at different sites, and at different times can be accurately removed, while preserving the intrinsic structure of the data.

Our main contributions are as follows. (i) We present a new implementation of PA using the Riemannian geometry of the Lorentz model for unsupervised label-free hierarchical data alignment. (ii) We provide theoretical analysis and justification of our alignment method based on new derivations of Riemannian geometry operations in the Lorentz model. These derivations have their own merit as they could be used in other contexts. (iii) We show experimental results of accurate batch effect removal from several hierarchical bioinformatics datasets without landmark correspondence.

2 Background on hyperbolic geometry

Hyperbolic space is a non-Euclidean space with a negative constant sectional curvature and an underlying geometry that describes tree-like graphs with small distortions [46]. There exist four commonly-used models for hyperbolic spaces: Poincaré disk model, Lorentz model (hyperboloid model), Poincaré half-plane model, and Beltrami-Klein model. These four models are equivalent and there exist transformations between them. Here, we consider the Lorentz model, and specifically, the upper sheet of the hyperboloid model, because its basic Riemannian operations have simple closed-form expressions and the computation of the geodesic distances is stable [42, 30].

Formally, the upper sheet of the hyperboloid model in a d -dimensional hyperbolic space is defined by $\mathbb{L}^d := \{\mathbf{x} \in \mathbb{R}^{d+1} | \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = -1, \mathbf{x}(1) > 0\}$, where $\langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = \mathbf{x}^{\top} \mathbf{H} \mathbf{x}$ is the Lorentzian inner product and $\mathbf{H} \in \mathbb{R}^{(d+1) \times (d+1)}$ is defined by $\mathbf{H} = [-1, \mathbf{0}^{\top}; \mathbf{0}, \mathbf{I}_d]$. The Lorentzian norm of a hyperbolic vector $\mathbf{x} \in \mathbb{L}^d$ is denoted by $\|\mathbf{x}\|_{\mathcal{L}} = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}}}$, with the origin $\boldsymbol{\mu}_0 = [1, \mathbf{0}^{\top}]^{\top} \in \mathbb{L}^d$. Let $\mathcal{T}_{\mathbf{x}}\mathbb{L}^d$ be the tangent space at $\mathbf{x} \in \mathbb{L}^d$, defined by $\mathcal{T}_{\mathbf{x}}\mathbb{L}^d := \{\mathbf{v} | \langle \mathbf{x}, \mathbf{v} \rangle_{\mathcal{L}} = 0\}$. Consider $\mathbf{x} \in \mathbb{L}^d$ and $\mathbf{v} \in \mathcal{T}_{\mathbf{x}}\mathbb{L}^d$, the geodesic path $\phi : \mathbb{R}_0^+ \rightarrow \mathbb{L}^d$ is defined by $\phi(t) = \cosh(\|\mathbf{v}\|_{\mathcal{L}}t)\mathbf{x} + \sinh(\|\mathbf{v}\|_{\mathcal{L}}t)\frac{\mathbf{v}}{\|\mathbf{v}\|_{\mathcal{L}}}$ with $\phi(0) = \mathbf{x}$ and initial velocity $\phi'(0) = \mathbf{v}$, where $\phi'(t) := \frac{d}{dt}\phi(t)$. In addition, the associated geodesic distance is $d_{\mathbb{L}^d}(\mathbf{x}, \phi_v(t)) = \cosh^{-1}(-\langle \mathbf{x}, \phi_v(t) \rangle_{\mathcal{L}})$.

The Exponential map, projecting a point $\mathbf{v} \in \mathcal{T}_{\mathbf{x}}\mathbb{L}^d$ to the manifold \mathbb{L}^d , is given by $\text{Exp}_{\mathbf{x}}(\mathbf{v}) = \phi(1) = \cosh(\|\mathbf{v}\|_{\mathcal{L}})\mathbf{x} + \sinh(\|\mathbf{v}\|_{\mathcal{L}})\frac{\mathbf{v}}{\|\mathbf{v}\|_{\mathcal{L}}}$. The Logarithmic map, projecting a point $\mathbf{y} \in \mathbb{L}^d$ to the tangent space $\mathcal{T}_{\mathbf{x}}\mathbb{L}^d$ at \mathbf{x} , is defined by $\text{Log}_{\mathbf{x}}(\mathbf{y}) = \frac{\cosh^{-1}(\lambda)}{\sqrt{\lambda^2 - 1}}(\mathbf{y} - \lambda\mathbf{x})$, where $\lambda = -\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}$. The PT of a vector $\mathbf{v} \in \mathcal{T}_{\mathbf{x}}\mathbb{L}^d$ along the geodesic path from $\mathbf{x} \in \mathbb{L}^d$ to $\mathbf{y} \in \mathbb{L}^d$ is defined by $\text{PT}_{\mathbf{x} \rightarrow \mathbf{y}}(\mathbf{v}) = \mathbf{v} + \frac{\langle \mathbf{y} - \lambda\mathbf{x}, \mathbf{v} \rangle_{\mathcal{L}}}{\lambda + 1}(\mathbf{x} + \mathbf{y})$, where $\lambda = -\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}$, keeping the metric tensor unchanged. The Riemannian mean $\bar{\mathbf{x}}_{\mathcal{X}}$ and the corresponding dispersion $d_{\mathcal{X}}$ of a set $\mathcal{X} = \{\mathbf{x}_i | \mathbf{x}_i \in \mathbb{L}^d\}_{i=1}^n$ are defined using the Fréchet mean [13, 33] by

$$\bar{\mathbf{x}}_{\mathcal{X}} := m(\mathcal{X}) = \arg \min_{\mathbf{x} \in \mathbb{L}^d} \sum_{i=1}^n d_{\mathbb{L}^d}^2(\mathbf{x}, \mathbf{x}_i) \text{ and } d_{\mathcal{X}} := r(\mathcal{X}) = \min_{\mathbf{x} \in \mathbb{L}^d} \frac{1}{n} \sum_{i=1}^n d_{\mathbb{L}^d}^2(\mathbf{x}, \mathbf{x}_i), \quad (1)$$




Figure 1: Illustration of our alignment method (HPA). Two sets of points in hyperbolic space are given, depicted in dark and light colors. Each point is associated with one of two labels: (i) blue/cyan, and (ii) red/pink. Left: Alignment results after each step of our HPA implemented in \mathbb{L}^2 : Riemannian translation, Riemannian scaling, and Riemannian wrapped rotation. Note that these Riemannian operations in hyperbolic space are different than their Euclidean counterparts. Right: Alignment results of four methods: PAH [51] and HOT-F [54] applied in \mathbb{L}^2 , and HOT-L [22] and HOT-ME [22] applied in the 2D Poincaré disk. For visualization, all points in \mathbb{L}^2 are transformed to the 2D Poincaré disk. The alignment after Step 3 of HPA (circled in black) is more accurate than the alignments obtained by the other methods.

where $m : \mathcal{X} \rightarrow \mathbb{L}^d$ and $r : \mathcal{X} \rightarrow \mathbb{R}^+$. Note that the Fréchet mean of samples on connected and compact Riemannian manifolds of non-positive curvatures, such as hyperbolic spaces, is guaranteed to exist, and it is unique [24, 44, 1]. The Fréchet mean is commonly computed by the Karcher Flow [24, 20], which is computationally demanding. Importantly, in the considered hyperbolic space, the Fréchet mean can be efficiently obtained using the accurate gradient formulation [33].

Given a vector $\bar{x} \in \mathbb{L}^d$ and a symmetric and positive-definite (SPD) matrix $\Sigma \in \mathbb{R}^{d \times d}$, the wrapped normal distribution $\mathcal{G}(\bar{x}, \Sigma)$ provides a generative model of hyperbolic samples as follows [39, 11]. First, a vector v' is sampled from $\mathcal{N}(\mathbf{0}, \Sigma)$. Then, 0 is concatenated to the vector v' such that $v = [0, v']^\top \in \mathcal{T}_{\mu_0} \mathbb{L}^d$. Finally, PT from the origin $\mu_0 = [1, \mathbf{0}^\top]^\top$ to \bar{x} is applied to v , and the resulting point is mapped to the manifold using the Exponential map at \bar{x} . The probability density function of this model is given by $\log \mathcal{G}(y|\bar{x}, \Sigma) = \log \mathcal{N}(\mathbf{0}, \Sigma) - (n-1) \log\left(\frac{\sinh\|v'\|_2}{\|v'\|_2}\right)$.

3 Hyperbolic Procrustes analysis

Existing methods for data alignment typically seek a function that minimizes a certain cost. A large body of work attempts to match the empirical densities of two datasets, e.g., by minimizing the maximum mean discrepancy (MMD) [48, 29] or solving OT problems [45, 2, 22]. Finding an effective cost function without labels or landmarks is challenging, and minimizing such costs directly often lead to poor alignment in practice (see illustration in Fig. 1). A different well-established approach that applies indirect alignment based on geometric considerations is PA. While preparing this manuscript, another method of PA in hyperbolic spaces (PAH) was presented for matching two sets, assuming that they consist of the same number of points and that there exists a point-wise isometric map between them [51]. We remark that the analysis we present here applies to broader settings and makes no such assumptions. See Appendix E for details on classical PA as well as for comparisons to [51] and to the application of Euclidean PA in the tangent space.

We consider two sets of points $\mathcal{H}^{(1)} = \{h_i^{(1)}\}_{i=1}^{N_1}$ and $\mathcal{H}^{(2)} = \{h_i^{(2)}\}_{i=1}^{N_2}$ in \mathbb{L}^d . Here, we aim to find a function $\zeta : \mathbb{L}^d \rightarrow \mathbb{L}^d$, consisting of three components: translation, scaling, and rotation, that aligns $\mathcal{H}^{(2)}$ with $\mathcal{H}^{(1)}$ in an unsupervised label-free manner as depicted in Fig. 1. Finding such a function can be viewed as an extension of classical PA from the Euclidean space \mathbb{R}^{d+1} to the Lorentz model \mathbb{L}^d . A natural extension to multiple sets is described in Section 3.5. We remark that the statements are written in the context of the problem at hand. In Appendix A we restate them more generally and present their proofs.

3.1 Riemannian translation

Let $\bar{h}^{(1)}$ and $\bar{h}^{(2)}$ denote the Riemannian means of the sets $\mathcal{H}^{(1)}$ and $\mathcal{H}^{(2)}$, respectively. In this translation component, we find a map $\Gamma_{\bar{h}^{(2)} \rightarrow \bar{h}^{(1)}} : \mathbb{L}^d \rightarrow \mathbb{L}^d$ that aligns the Riemannian means of

the sets. In the spirit of [5, 53], we propose to construct $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)})$ as the composition of three Riemannian operations in \mathbb{L}^d : the Logarithmic map applied to $\mathbf{h}_i^{(2)}$ at $\bar{\mathbf{h}}^{(2)}$, PT $\mathbf{h}_i^{(2)}$ from $\bar{\mathbf{h}}^{(2)}$ to $\bar{\mathbf{h}}^{(1)}$ along the geodesic path, and the Exponential map applied to the transported point at $\bar{\mathbf{h}}^{(1)}$:

$$\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)}) := \text{Exp}_{\bar{\mathbf{h}}^{(1)}}(\text{PT}_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\text{Log}_{\bar{\mathbf{h}}^{(2)}}(\mathbf{h}_i^{(2)}))). \quad (2)$$

See Fig. B.1 in Appendix B for illustration. Since the geodesic path between any two points in \mathbb{L}^d is unique [46], $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$ is well-defined. The rationale behind the combination of these three Riemannian operations is twofold. First, PT is a map that aligns the means of the sets, while preserving their internal structure. Second, the Logarithmic and Exponential maps compose a map whose domain and range are the Lorentz model \mathbb{L}^d rather than the tangent space, as desired. We make these claims formal in the following results.

Proposition 1. *The map $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$ defined in Eq. (2) aligns the means of the sets, i.e., it satisfies*

$$\bar{\mathbf{h}}^{(1)} = m(\{\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)})\}_{i=1}^{N_2}), \quad (3)$$

where m is the function defined in Eq. (I).

Proposition 2. *The map $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)})$ for all $\mathbf{h}_i^{(2)} \in \mathcal{H}^{(2)}$ can be recast as:*

$$\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)}) = \mathbf{h}_i^{(2)} - \beta(\mathbf{h}_i^{(2)} | \bar{\mathbf{h}}^{(1)}, \bar{\mathbf{h}}^{(2)}) \bar{\mathbf{h}}^{(2)} + \gamma(\mathbf{h}_i^{(2)} | \bar{\mathbf{h}}^{(1)}, \bar{\mathbf{h}}^{(2)}) \bar{\mathbf{h}}^{(1)}, \quad (4)$$

where the functions β and γ are positive, defined by $0 < \beta(\mathbf{h}_i^{(2)} | \bar{\mathbf{h}}^{(1)}, \bar{\mathbf{h}}^{(2)}) = -\left\langle \frac{\bar{\mathbf{h}}^{(1)} + \bar{\mathbf{h}}^{(2)}}{\alpha + 1}, \mathbf{h}_i^{(2)} \right\rangle_{\mathcal{L}}$ and $0 < \gamma(\mathbf{h}_i^{(2)} | \bar{\mathbf{h}}^{(1)}, \bar{\mathbf{h}}^{(2)}) = \left\langle \frac{\bar{\mathbf{h}}^{(1)} - (2\alpha + 1)\bar{\mathbf{h}}^{(2)}}{\alpha + 1}, \mathbf{h}_i^{(2)} \right\rangle_{\mathcal{L}}$, respectively, and $0 < \alpha = -\langle \bar{\mathbf{h}}^{(1)}, \bar{\mathbf{h}}^{(2)} \rangle_{\mathcal{L}}$.

In addition to providing a compact closed-form expression, Prop. 2 gives the proposed translation based on Riemannian geometry an interpretation of standard mean alignment in linear vector spaces. It implies that the alignment is nothing but subtracting the mean of the source set $\bar{\mathbf{h}}^{(2)}$ from each vector in $\mathcal{H}^{(2)}$, and adding the mean of the target set $\bar{\mathbf{h}}^{(1)}$ (with the appropriate scales).

Proposition 3. *The map $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$ preserves distances (i.e., it is an isometry):*

$$d_{\mathbb{L}^d}(\mathbf{h}_i^{(2)}, \mathbf{h}_j^{(2)}) = d_{\mathbb{L}^d}(\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(2)}), \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{h}_j^{(2)})), \quad (5)$$

for any two points $\mathbf{h}_i^{(2)}, \mathbf{h}_j^{(2)} \in \mathcal{H}^{(2)}$.

Let $\psi(t)$ be the unique geodesic path from $\bar{\mathbf{h}}^{(2)}$ to $\bar{\mathbf{h}}^{(1)}$ such that $\psi(0) = \bar{\mathbf{h}}^{(2)}$ and $\psi(1) = \bar{\mathbf{h}}^{(1)}$, and let $\psi'(0) \in \mathcal{T}_{\bar{\mathbf{h}}^{(2)}}\mathbb{L}^d$ and $\psi'(1) \in \mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d$ be the corresponding velocities, respectively.

Proposition 4. *The map $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$ aligns geodesic velocities, i.e., given the mapping of the geodesic velocities to the manifold $\mathbb{L}^d \ni \mathbf{v}_0 = \text{Exp}_{\bar{\mathbf{h}}^{(2)}}(\psi'(0)) = \bar{\mathbf{h}}^{(1)}$ and $\mathbb{L}^d \ni \mathbf{v}_1 = \text{Exp}_{\bar{\mathbf{h}}^{(1)}}(\psi'(1))$, we have*

$$\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathbf{v}_0) = \mathbf{v}_1. \quad (6)$$

Isometry is determined up to rotation, a fact that can be problematic for alignment. For example, any H-unitary matrix [16] can be an isometric function in \mathbb{L}^d . When landmarks are given, they can be used to alleviate this redundancy. However, in the purely unsupervised setting we consider, other data-driven cues are required. Prop. 4 implies that the proposed translation based on PT fixes some of these rotational degrees of freedom by aligning the geodesic velocities. In Section 3.3 we revisit this issue. Now, with a slight abuse of notation, let $\tilde{\mathcal{H}}^{(2)} = \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathcal{H}^{(2)})$.

Proposition 5. *Consider two subsets $\mathcal{A}, \mathcal{B} \subset \mathcal{H}^{(2)}$ and their translations $\tilde{\mathcal{A}} = \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathcal{A})$, $\tilde{\mathcal{B}} = \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}(\mathcal{B}) \subset \tilde{\mathcal{H}}^{(2)}$. Let $\bar{\mathbf{a}} = m(\mathcal{A})$, $\bar{\mathbf{b}} = m(\mathcal{B})$, $\tilde{\mathbf{a}} = m(\tilde{\mathcal{A}})$, and $\tilde{\mathbf{b}} = m(\tilde{\mathcal{B}})$ be the Riemannian means of the subsets. Then,*

$$\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}} \circ \Gamma_{\bar{\mathbf{a}} \rightarrow \bar{\mathbf{b}}} = \Gamma_{\tilde{\mathbf{a}} \rightarrow \tilde{\mathbf{b}}} \circ \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}. \quad (7)$$

In the context of the alignment problem, the importance of Prop. 5 is the following. Suppose the two sets correspond to data measured at two labs (denoted with and without a tilde), and suppose each set was acquired by two types of equipment (denoted by \mathcal{A} and \mathcal{B}). Prop. 5 implies that aligning data from the different labs and then aligning data acquired using the different equipment is equivalent to first aligning the different equipment and then the different labs, i.e., any order of the two alignments generates the same result. Seemingly, this is a natural property in Euclidean spaces. However, in a Riemannian manifold, it is not a trivial result, and it holds for the transport along the geodesic path. See Appendix A for counter-examples.

3.2 Riemannian scaling

Let $d^{(1)}$ and $d^{(2)}$ denote the Riemannian dispersions of $\mathcal{H}^{(1)}$ and $\mathcal{H}^{(2)}$. By propositions 1 and 3 $\bar{\mathbf{h}}^{(1)}$ and $d^{(2)}$ are the mean and dispersion of $\tilde{\mathcal{H}}^{(2)}$. Here, our goal is to align the Riemannian dispersions of $\mathcal{H}^{(1)}$ and $\tilde{\mathcal{H}}^{(2)}$. For this purpose, we propose the scaling function $\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s : \mathbb{L}^d \rightarrow \mathbb{L}^d$, given by

$$\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s(\tilde{\mathbf{h}}_i^{(2)}) = \phi_i(s), \quad (8)$$

where $s = \sqrt{d^{(1)}/d^{(2)}}$ is the scaling factor and $\phi_i(t)$ is the geodesic path from $\bar{\mathbf{h}}^{(1)}$ to $\tilde{\mathbf{h}}_i^{(2)}$ such that $\phi_i(0) = \bar{\mathbf{h}}^{(1)}$ and $\phi_i(1) = \tilde{\mathbf{h}}_i^{(2)}$. See Fig. B.2 in Appendix B for illustration.

Proposition 6. *The dispersion of the rescaled set $\hat{\mathcal{H}}^{(2)} = \Upsilon_{\bar{\mathbf{h}}^{(1)}}^s(\tilde{\mathcal{H}}^{(2)})$ is $d^{(1)}$.*

3.3 Riemannian wrapped rotation

The purpose of this component is to align the orientation of the distributions of the two sets after translation and scaling, namely, after aligning their first and second moments. The proposed rotation function $\Theta_{\bar{\mathbf{h}}^{(1)}} : \mathbb{L}^d \rightarrow \mathbb{L}^d$ consists of (i) mapping the points from the manifold \mathbb{L}^d to the tangent space $\mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d$, (ii) mapping to \mathbb{R}^d , (iii) rotating in \mathbb{R}^d , and (iv) mapping back to the tangent space and then to the manifold. We perform the rotation in \mathbb{R}^d , which we term *wrapped rotation*, rather than a direct rotation on the manifold \mathbb{L}^d or on the tangent space $\mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d$ for the following reasons. First, the frequently used rotation map in \mathbb{L}^d [5] does not necessarily preserve the Riemannian mean, and in our context, it might reverse the mean alignment. Second, rotation applied directly to the tangent space $\mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d$ does not guarantee that the rotated points remain on the same tangent space. Third, applying rotation in the Euclidean vector space that is isometric to the tangent space is less efficient and stable, and it obtains slightly worse empirical results (see Appendix D.4 for details). Last, applying the rotation in \mathbb{R}^d allows us to use the standard Euclidean rotation using SVD. In Section 4, we empirically demonstrate the advantage of the proposed rotation compared to the alternatives.

Definition 1. *Let the mapping function $\mathcal{P}_{\bar{\mathbf{h}}^{(1)}} : \mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d \rightarrow \mathbb{R}^d$ defined on the tangent space at $\bar{\mathbf{h}}^{(1)}$ and its inverse map be the following functions defined by*

$$\mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\mathbf{v}) := [\mathbf{v}(2), \dots, \mathbf{v}(d+1)]^\top \in \mathbb{R}^d \text{ and } \mathcal{P}_{\bar{\mathbf{h}}^{(1)}}^{-1}(\mathbf{s}) := \left[\frac{\langle \mathbf{s}, \mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\bar{\mathbf{h}}^{(1)}) \rangle}{\bar{\mathbf{h}}^{(1)}(1)}, \mathbf{s}^\top \right]^\top \in \mathcal{T}_{\bar{\mathbf{h}}^{(1)}}\mathbb{L}^d, \quad (9)$$

where $\mathbf{s} \in \mathbb{R}^d$ and $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product.

Note that removing the first element of \mathbf{v} is valid due to the constraint imposed on the vector elements in the tangent space by definition. Indeed, no information is lost and the mapping is invertible.

The first step in our rotation component is to map the points in $\mathcal{H}^{(1)}$ and in $\hat{\mathcal{H}}^{(2)}$ to the tangent space at $\bar{\mathbf{h}}^{(1)}$: $\mathbf{v}_i^{(1)} = \text{Log}_{\bar{\mathbf{h}}^{(1)}}(\mathbf{h}_i^{(1)})$ for $i = 1, \dots, N_1$, and $\mathbf{v}_i^{(2)} = \text{Log}_{\bar{\mathbf{h}}^{(1)}}(\hat{\mathbf{h}}_i^{(2)})$ for $i = 1, \dots, N_2$. In the second step, we map the points by the mapping function in Definition 1 and re-center them: $\mathbf{s}_i^{(1)} = \mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\mathbf{v}_i^{(1)}) - \bar{\mathbf{s}}^{(1)}$, $i = 1, \dots, N_1$ and $\mathbf{s}_i^{(2)} = \mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\mathbf{v}_i^{(2)}) - \bar{\mathbf{s}}^{(2)}$, $i = 1, \dots, N_2$, where $\bar{\mathbf{s}}^{(k)} = \frac{1}{N_k} \sum_{i=1}^{N_k} \mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\mathbf{v}_i^{(k)})$ for $k = 1, 2$ is the mean vector of the projections. Then, the mapped and centered points (in \mathbb{R}^d) are collected into matrices:

$$\mathbf{S}^{(k)} = [\mathbf{s}_1^{(k)}, \mathbf{s}_2^{(k)}, \dots, \mathbf{s}_{N_k}^{(k)}] \in \mathbb{R}^{d \times N_k}. \quad (10)$$

In the third step, for each set $k = 1, 2$, we compute the rotation matrix $\mathbf{U}^{(k)} \in \mathbb{R}^{d \times d}$ by applying SVD to the matrix $\mathbf{S}^{(k)} = \mathbf{U}^{(k)} \mathbf{\Lambda}^{(k)} (\mathbf{E}^{(k)})^\top$. Since the left-singular vectors are determined up to a sign, we propose to align their signs as follows: $\mathbf{u}_i^{(2)} \leftarrow \text{sign}(\langle \mathbf{u}_i^{(2)}, \mathbf{u}_i^{(1)} \rangle) \mathbf{u}_i^{(2)}$, where $\mathbf{u}_i^{(1)}$ and $\mathbf{u}_i^{(2)}$ are the i -th left-singular vector of the two sets, resulting in modified rotation matrices $\mathbf{U}^{(k)}$. Finally, we apply the rotation to $\widehat{\mathcal{H}}^{(2)}$ by

$$\Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}}(\widehat{\mathbf{h}}_i^{(2)}) = \text{Exp}_{\bar{\mathbf{h}}^{(1)}} \left(\mathcal{P}_{\bar{\mathbf{h}}^{(1)}}^{-1} \left(\mathbf{U}^\top \left(\mathcal{P}_{\bar{\mathbf{h}}^{(1)}}(\text{Log}_{\bar{\mathbf{h}}^{(1)}}(\widehat{\mathbf{h}}_i^{(2)})) - \bar{\mathbf{s}}^{(2)} \right) + \bar{\mathbf{s}}^{(2)} \right) \right), \quad (11)$$

where $\mathbf{U} = \mathbf{U}^{(1)}(\mathbf{U}^{(2)})^\top$.

Proposition 7. *The wrapped rotation is bijective, and the inverse is given by*

$$(\Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}})^{-1} = \Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}^\top}. \quad (12)$$

3.4 Analysis

Putting all three components together, the proposed HPA that aligns $\mathcal{H}^{(2)}$ with $\mathcal{H}^{(1)}$ culminates in the composition of translation, scaling, and rotation:

$$\Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}} \circ \Upsilon_{\bar{\mathbf{h}}^{(1)}}^s \circ \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}. \quad (13)$$

As in most PA schemes, the order of the three components is important. Yet, the proposed components allow us a certain degree of freedom, as indicated in the following results.

Proposition 8. *The Riemannian translation and the Riemannian scaling commute w.r.t. the Riemannian means $\bar{\mathbf{h}}^{(1)}$ and $\bar{\mathbf{h}}^{(2)}$:*

$$\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s \circ \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}} = \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}} \circ \Upsilon_{\bar{\mathbf{h}}^{(2)}}^s. \quad (14)$$

Note that $\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s$ and $\Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$ do not necessarily commute: $\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s \circ \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}} \neq \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}} \circ \Upsilon_{\bar{\mathbf{h}}^{(1)}}^s$.

Proposition 9. *The Riemannian scaling and the wrapped rotation commute:*

$$\Upsilon_{\bar{\mathbf{h}}^{(1)}}^s \circ \Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}} = \Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}} \circ \Upsilon_{\bar{\mathbf{h}}^{(1)}}^s. \quad (15)$$

We note that the rotation does not commute with the translation, because PT only preserves the local covariant derivative on the tangent space but might cause rotation and distortion along the transportation. Therefore, the rotation is required to be the last component of our HPA.

Thus far, we did not present a model for the discrepancy between the two sets, nor we presented the proposed HPA as optimal with respect to some criterion. In the following result, we show that if the discrepancy between the sets can be expressed as a composition of translation, scaling, and rotation, then the two sets can be perfectly aligned using HPA.

Proposition 10. *Let $\eta : \mathbb{L}^d \rightarrow \mathbb{L}^d$ be a map, given by $\eta = \Theta_{\bar{\mathbf{h}}^{(1)}}^{\mathbf{U}} \circ \Upsilon_{\bar{\mathbf{h}}^{(1)}}^s \circ \Gamma_{\bar{\mathbf{h}}^{(2)} \rightarrow \bar{\mathbf{h}}^{(1)}}$. If $\mathcal{H}^{(1)} = \{\mathbf{h}_i^{(1)} = \eta(\mathbf{h}_i^{(2)})\}_{i=1}^{N_2}$, then,*

$$\mathbf{h}_i^{(2)} = (\Theta_{\bar{\mathbf{h}}^{(2)}}^{\mathbf{U}'} \circ \Upsilon_{\bar{\mathbf{h}}^{(2)}}^{\frac{1}{s}} \circ \Gamma_{\bar{\mathbf{h}}^{(1)} \rightarrow \bar{\mathbf{h}}^{(2)}})(\mathbf{h}_i^{(1)}), \quad i = 1, \dots, N_2, \quad (16)$$

where $\mathbf{U}' \in \mathbb{O}(d)$.

Note that HPA consists of the sequence of Riemannian translation, Riemannian scaling and wrapped-rotation. The domain and range of each component is the manifold \mathbb{L}^d . Yet, the first and last operations of each component are the Logarithmic and Exponential maps that project a point from the manifold to the tangent space, and vice versa, respectively. This allows us to propose an efficient implementation of the sequence without the back and forth projections as described in Appendix C

3.5 Extension to multiple sets

We can naturally scale up the setting to support the alignment of $K > 2$ sets, denoted by $\mathcal{H}^{(k)} = \{\mathbf{h}_i^{(k)}\}_{i=1}^{N_k}$, where $k \in \{1, 2, \dots, K\}$. Let $\bar{\mathbf{h}}^{(k)}$ and $d^{(k)}$ be the Riemannian mean and dispersion of

Algorithm 1 Hyperbolic Procrustes analysis

Input: K sets of hyperbolic points $\mathcal{H}^{(1)} = \{\mathbf{h}_i^{(1)}\}_{i=1}^{N_1}, \dots, \mathcal{H}^{(K)} = \{\mathbf{h}_i^{(K)}\}_{i=1}^{N_K}$

Output: K aligned sets of hyperbolic points $\check{\mathcal{H}}^{(1)} = \{\check{\mathbf{h}}_i^{(1)}\}_{i=1}^{N_1}, \dots, \check{\mathcal{H}}^{(K)} = \{\check{\mathbf{h}}_i^{(K)}\}_{i=1}^{N_K}$

- 1: **for** each set $\mathcal{H}^{(k)}$ **do**
 - 2: compute the Riemannian mean $\bar{\mathbf{h}}^{(k)}$ and dispersion $d^{(k)}$
 - 3: **end for**
 - 4: compute $\bar{\mathbf{h}}$, the global Riemannian mean of $\{\bar{\mathbf{h}}^{(k)}\}_{k=1}^K$
 - 5: **for** each set $\mathcal{H}^{(k)}$ **do**
 - 6: apply the Riemannian translation $\tilde{\mathbf{h}}_i^{(k)} = \Gamma_{\bar{\mathbf{h}}^{(k)} \rightarrow \bar{\mathbf{h}}}(\mathbf{h}_i^{(k)})$ // Eq. (2)
 - 7: apply the Riemannian scaling $\hat{\mathbf{h}}_i^{(k)} = \Upsilon_{\bar{\mathbf{h}}}^s(\tilde{\mathbf{h}}_i^{(k)})$ with $s = 1/\sqrt{d^{(k)}}$ // Eq. (8)
 - 8: apply the wrapped rotation $\check{\mathbf{h}}_i^{(k)} = \Theta_{\bar{\mathbf{h}}}^{\mathbf{U}}(\hat{\mathbf{h}}_i^{(k)})$ with $\mathbf{U} = \mathbf{U}^{(1)}(\mathbf{U}^{(k)})^\top$ // Eq. (11)
 - 9: **end for**
-

the k -th dataset, respectively. In addition, let $\bar{\mathbf{h}}$ be the global Riemannian mean of $\{\bar{\mathbf{h}}^{(k)}\}_{k=1}^K$. We propose to transport the points of the k -th set using $\Gamma_{\bar{\mathbf{h}}^{(k)} \rightarrow \bar{\mathbf{h}}}$. Next, the Riemannian dispersion of each set is set to 1 by applying $\Upsilon_{\bar{\mathbf{h}}}^s$ with $s = 1/\sqrt{d^{(k)}}$. Finally, the wrapped rotation is applied to all the data sets on the mapping of the tangent space $\mathcal{T}_{\bar{\mathbf{h}}}\mathbb{L}^d$ and then mapped back to the manifold \mathbb{L}^d . The first set is designated as the reference set, and all other rotation matrices $\mathbf{U}^{(k)}$ are updated according to $\mathbf{u}_i^{(k)} \leftarrow \text{sign}(\langle \mathbf{u}_i^{(k)}, \mathbf{u}_i^{(1)} \rangle) \mathbf{u}_i^{(k)}$. The proposed HPA for multiple sets is summarized in Algorithm 1, and some implementation remarks appear in Appendix C.

4 Experimental results

We apply HPA to simulations and to three biomedical datasets [1]. In addition, we test HPA on MNIST [31] and USPS [23] datasets, which arguably do not have a distinct hierarchical structure. Nonetheless, we demonstrate in Appendix D that our HPA is highly effective in aligning these two datasets. All the experiments are label-free. We compare the obtained results to the following alignment methods: (i) PAH [51], which is applied only to the simulated data since it requires the existence of a one-to-one correspondence between the sets, (ii) only the Riemannian translation (RT), (iii) OT in hyperbolic space with the weighted Fréchet mean (HOT-F) extended to an unsupervised setting according to [54], (iv) OT with W-linear map (HOT-L) [22], and (v) hyperbolic mapping estimation (HOT-ME) [22]. As a baseline, we present the results obtained before the alignment (Baseline). For more details on the experimental setting, see Appendix C.

4.1 Simulations

The synthetic data in \mathbb{L}^d is generated using the sampling scheme described in Section 2 based on [39]. Given an arbitrary point $\boldsymbol{\mu} \in \mathbb{L}^d$ and an arbitrary SPD matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$, we generate a set of N points $\mathcal{Q}^{(1)} = \{\mathbf{q}_i^{(1)}\}_{i=1}^N$ centered at $\boldsymbol{\mu}$ by $\mathbb{L}^d \ni \mathbf{q}_i^{(1)} = \text{Exp}_{\boldsymbol{\mu}}(\text{PT}_{\boldsymbol{\mu}_0 \rightarrow \boldsymbol{\mu}}(\tilde{\mathbf{v}}_i^{(1)}))$, where $\boldsymbol{\mu}_0 = [1, \mathbf{0}]^\top$ is the origin, $\mathbf{v}_i^{(1)} = [0, \tilde{\mathbf{v}}_i^{(1)}]^\top$, and $\tilde{\mathbf{v}}_i^{(1)} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. Next, we generate three noisy and distorted versions of $\mathcal{Q}^{(1)}$. The first noisy set $\mathcal{Q}^{(2)} = \{\mathbf{q}_i^{(2)}\}_{i=1}^N$ is generated as proposed in [51] by $\mathbf{q}_i^{(2)} = \mathbf{L} \mathbf{T}_{\boldsymbol{\epsilon}_i} \mathbf{q}_i^{(1)}$, where $\mathbf{T}_{\boldsymbol{\epsilon}_i}$ is a hyperbolic translation defined by $\mathbf{T}_{\boldsymbol{\epsilon}_i} = [\sqrt{1 + \boldsymbol{\epsilon}_i^\top \boldsymbol{\epsilon}_i}, \boldsymbol{\epsilon}_i^\top; \boldsymbol{\epsilon}_i, (\mathbf{I} + \boldsymbol{\epsilon}_i \boldsymbol{\epsilon}_i^\top)^{\frac{1}{2}}]$, $\boldsymbol{\epsilon}_i$ is sampled from $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, σ^2 is the variance, and \mathbf{L} is a random H-unitary matrix [16]. Another noisy set, denoted as $\mathcal{Q}^{(3)} = \{\mathbf{q}_i^{(3)}\}_{i=1}^N$, is generated by $\mathbf{q}_i^{(3)} = \mathbf{L}(\text{Exp}_{\boldsymbol{\mu}}(\text{PT}_{\boldsymbol{\mu}_0 \rightarrow \boldsymbol{\mu}}(\mathbf{u}_i^{(1)})))$, where $\mathbf{u}_i^{(1)} = [0, \tilde{\mathbf{v}}_i^{(1)} + \boldsymbol{\epsilon}_i]^\top$. Here, the noise is added to the tangent space at $\boldsymbol{\mu}_0$. Finally, let $\mathcal{Q}^{(4)} = \{\mathbf{q}_i^{(4)}\}_{i=1}^N$ be a distorted set, given by $\mathbf{q}_i^{(4)} = f_{\boldsymbol{\mu}'}(\mathbf{q}_i^{(3)})$, where $f_{\boldsymbol{\mu}'}(\mathbf{x}) = \cosh(\|\mathbf{u}\|_{\mathcal{L}t})\boldsymbol{\mu}' + \sinh(\|\mathbf{u}\|_{\mathcal{L}t})\frac{\mathbf{u}}{\|\mathbf{u}\|_{\mathcal{L}}}$ and $\mathbf{u} = \text{Log}_{\boldsymbol{\mu}'}(\mathbf{x})$, for arbitrary (fixed) $\boldsymbol{\mu}' \in \mathbb{L}^d$ and $t > 0$.

¹Our code is available at <https://github.com/RonenTalmonLab/HyperbolicProcrustesAnalysis>




Figure 2: Alignment results of the simulated data.




Figure 3: The visualizations of HPA applied to the two breast cancer gene expression datasets.

We apply Algorithm 1 to align the three pairs of sets $\{Q^{(1)}, Q^{(2)}\}$, $\{Q^{(1)}, Q^{(3)}\}$, and $\{Q^{(1)}, Q^{(4)}\}$, setting $N = 100$, $\sigma = 1$, and $d \in \{3, 5, 10, 20, \dots, 40\}$. Each experiment is repeated 10 times with different values of μ , Σ , μ' and t . To evaluate the alignment, we use the pairwise discrepancy based on the *hidden* one-to-one correspondence, given by $\varepsilon(Q^{(1)}, Q^{(j)}) = \frac{1}{N} \sum_{i=1}^N d_{\mathbb{L}^d}^2(q_i^{(1)}, q_i^{(j)})$, where $j \in \{2, 3, 4\}$. The discrepancy as a function of the dimension d is shown in Fig. 2. We observe that the proposed HPA has lower discrepancy relative to the other label-free methods. Specifically, it outperforms OT-based methods that are designed to match the densities. Furthermore, the proposed HPA is stable, in contrast to HOT-L, which is highly sensitive to the noise and distortion introduced in $Q^{(3)}$ and $Q^{(4)}$. Interestingly, we remark that the discrepancies of RT and PAH are very close, empirically showing that RT alone is comparable to PAH. In addition, note that HPA is permutation-invariant and does not require one-to-one correspondence as PAH. We report the running-time in Appendix D and demonstrate that HPA is more efficient than HOT-F and HOT-ME.

4.2 Batch effect removal

We consider bioinformatics datasets consisting of gene expression data and CyTOF. Representing such data in hyperbolic spaces was shown to be informative and useful [25], implying that such data have an underlying inherent hierarchical structure. Batch effects [43] arise from experimental variations that can be attributed to the measurement device or other environmental factors. Batch correction is typically a critical precursor to any subsequent analysis and processing.

Three batch effect removal tasks are examined. The first task involves breast cancer (BC) gene expression data. We consider two publicly available datasets: METABRIC [8] and TCGA [26], consisting of samples from five breast cancer subtypes. The batch effect stems from different profiling techniques: gene expression microarray and RNA sequencing. In the second task, three cohorts of lung cancer (LC) gene expression data [21] are considered, consisting of samples from three lung cancer subtypes. The data were collected using gene expression microarrays at three different sites (a likely source of batch effects): Stanford University (ST), University of Michigan (UM), and Dana-Farber Cancer Institute (D-F). The last task involves CyTOF data [48] consisting of peripheral blood mononuclear cells (PBMCs) collected from two multiple sclerosis patients during four days: two day before treatment (BT) and two days after treatment (AT). These $8 = 2 \times 2 \times 2$ batches were collected with or without PMA/ionomycin stimulated PBMCs. We aim to remove the batch effects between two different days from the same condition (BT/AT) and from the same patient. In each batch removal task, we first learn an embedding of the data from all the batches into the Lorentz model \mathbb{L}^d [42]. Then, HPA is applied to the embedded points in \mathbb{L}^d .

Table 1: The k-NN AUC-ROC from the best k in each method.

Datasets	Space	S-Baseline	Baseline	HPA	RT	HOT-F	HOT-L	HOT-ME
BC	\mathbb{L}^{10}	0.7716 \pm 0.0088	0.5254 \pm 0.0091	0.7410 \pm 0.0354	0.6001 \pm 0.0658	0.5838 \pm 0.0121	0.4594 \pm 0.0428	0.5600 \pm 0.0147
LC	\mathbb{L}^{20}	0.9137 \pm 0.0532	0.5521 \pm 0.0991	0.8316 \pm 0.0904	0.5318 \pm 0.0370	0.5400 \pm 0.0396	0.5150 \pm 0.0654	0.4901 \pm 0.0516
P1 BT	\mathbb{L}^{20}	0.9723 \pm 0.0058	0.6646 \pm 0.1556	0.9401 \pm 0.0068	0.9020 \pm 0.0086	0.8603 \pm 0.0142	0.8855 \pm 0.0057	0.8919 \pm 0.0089
P1 AT	\mathbb{L}^{20}	0.9590 \pm 0.0150	0.7656 \pm 0.1564	0.9329 \pm 0.0011	0.8270 \pm 0.0880	0.8469 \pm 0.0873	0.8914 \pm 0.0175	0.8848 \pm 0.0306
P2 BT	\mathbb{L}^{20}	0.9686 \pm 0.0114	0.6971 \pm 0.1335	0.9329 \pm 0.0186	0.8830 \pm 0.0142	0.8762 \pm 0.0215	0.8566 \pm 0.1720	0.8810 \pm 0.0040
P2 AT	\mathbb{L}^{20}	0.9045 \pm 0.0070	0.5688 \pm 0.0688	0.8453 \pm 0.0798	0.7190 \pm 0.0439	0.7012 \pm 0.0912	0.7136 \pm 0.0012	0.7291 \pm 0.0113

Table 2: The mean and std of MMD computed over ten random subsets of size 300 for BC, 15 for LC, and 1000 for CyTOF from the two datasets.

Datasets	Space	Baseline	HPA	RT	HOT-F	HOT-L	HOT-ME
BC	\mathbb{L}^{10}	0.2089 \pm 0.0027	0.0013 \pm 0.0004	0.0072 \pm 0.0011	0.0009 \pm 0.0001	0.0019 \pm 0.0004	0.0007 \pm 0.0002
ST&UM	\mathbb{L}^{20}	0.1072 \pm 0.0051	0.0162 \pm 0.0048	0.0250 \pm 0.0049	0.0023 \pm 0.0002	0.0014 \pm 0.0007	0.0011 \pm 0.0006
ST&D-F	\mathbb{L}^{20}	0.3213 \pm 0.0152	0.0122 \pm 0.0042	0.0150 \pm 0.0106	0.0019 \pm 0.0002	0.0010 \pm 0.0004	0.0015 \pm 0.0003
UM&D-M	\mathbb{L}^{20}	0.0790 \pm 0.0071	0.0168 \pm 0.0090	0.0168 \pm 0.0032	0.0017 \pm 0.0003	0.0012 \pm 0.0006	0.0029 \pm 0.0005
P1 BT	\mathbb{L}^{20}	0.0638 \pm 0.0024	0.0012 \pm 0.0002	0.0020 \pm 0.0002	0.0004 \pm 0.0002	0.0002 \pm 0.0001	0.0009 \pm 0.0002
P1 AT	\mathbb{L}^{20}	0.0598 \pm 0.0014	0.0006 \pm 0.0001	0.0015 \pm 0.0001	0.0003 \pm 0.0001	0.0001 \pm 0.0001	0.0002 \pm 0.0001
P2 BT	\mathbb{L}^{20}	0.0424 \pm 0.0021	0.0012 \pm 0.0001	0.0015 \pm 0.0001	0.0002 \pm 0.0001	0.0020 \pm 0.0008	0.0009 \pm 0.0003
P2 AT	\mathbb{L}^{20}	0.0758 \pm 0.0053	0.0011 \pm 0.0002	0.0013 \pm 0.0002	0.0002 \pm 0.0001	0.0015 \pm 0.0008	0.0007 \pm 0.0003

Fig. 3 shows a visualization of the embedding of the two breast cancer datasets before and after HPA. For visualization, we project the points in \mathbb{L}^3 to the 3D Poincaré ball. Before the alignment, the dominant factor separating between the patients’ samples (points) is the batch. In contrast, after the alignment, the batch effect is substantially suppressed (visually) and the factors separating the points are dominated by the cancer subtype.

We evaluate the quality of the alignment in two aspects using objective measures: (i) k-NN classification, with leave-one-batch-out cross-validation, is utilized for assessing the alignment of the intrinsic structure, and (ii) MMD [19] is used for assessing the distribution alignment quality. For the classification, we view the five subtypes of BC, the three subtypes of LC, and the presence of stimulated cells in CyTOF as the labels in the respective tasks. In addition to the results of the different alignment methods, we report the k-NN classification based only on a single batch (*S-Baseline*), which indicates the adequacy of the representation in hyperbolic space to the task at hand.

Table 1 depicts the k-NN classification obtained for the best k per method, and Table 2 shows the MMD. In each task, we set the dimension of the Lorentz model d to the dimension in which the best empirical single-task performance is obtained (S-Baseline). We note that similar results and trends are obtained for various dimensions. Additional results for various k values and an ablation study, showing that the combination of all three components yields the best classification results, are reported in Appendix D.

Although the OT-based methods obtain the best matching between the distributions of the batches, HPA outperforms in all three tasks in terms of classification (see Table 1). In the two gene expression tasks, where the data have multiple labels and we align multiple batches, the advantage of HPA compared to the other methods is particularly significant. In Appendix D, we demonstrate HPA’s out-of-sample capabilities on the CyTOF data by learning the batch correction map between the different days from one patient and applying it to the data of the other patient.

4.3 Discussion

Alignment methods based on density matching, such as OT-based methods, often overlook an important aspect in purely unsupervised settings. Although sample density is the main data property that can be and need to be aligned, preserving the intrinsic structure/geometry of the sets is important, as it might be tightly related to the hidden labels. Indeed, we see in our experiments that OT-based methods provide a good density alignment (reducing the inter-set variability), as demonstrated by small MMD values (see Table 2). However, the intrinsic structure of the sets (the intra-set variability) is not preserved, as evident by the resulting poor (hidden) label matching, conveyed by the k-NN classification performance (see Table 1). This is also illustrated in the right panel of Fig. 3. There it is visible that the three OT-based methods provide good global alignment of the sets, yet the intrinsic structure is not kept, as implied by the poor color matching.

In contrast to OT-based methods, HPA does not explicitly aim to match densities, and thus, it obtains slightly worse MMD performance compared to OT-based methods. However, HPA matches the first two moments of the density and includes the rotation component, which was shown to be one of the fundamental limitations of OT-based methods for alignment as OT cannot recover volume-preserving maps [4, 36]. As seen in the simulation and experimental results and illustrated in the right panel of Fig. 1, we still obtain a good global alignment and simultaneously preserve the intrinsic structure, allowing for high classification performance. We remark that in the synthetic examples, there is a (hidden) one-to-one correspondence between the sets, and therefore one-to-one discrepancy can be computed (instead of or in addition to MMD). When there is such a correspondence, OT still cannot recover volume-preserving maps, while HPA can mitigate noise and distortions.

5 Conclusion

We introduced HPA for label-free alignment of data in the Lorentz model. Based on Riemannian geometry, we presented new translation and scaling operations that align the first and second Riemannian moments as well as a wrapped rotation that aligns the orientation in the hyperboloid model. Our theoretical analysis provides further insight and highlights properties that may be useful for practitioners. We empirically showed in simulations that HPA is stable under noise and distortions. Experimental results involving purely unsupervised batch correction of multiple bioinformatics datasets with multiple labels is demonstrated. Beyond alignment and batch effect removal, our method can be viewed as a type of domain adaptation or a precursor of transfer learning that relies on purely geometric considerations, exploiting the geometric structure of data as well as the geometric properties of the space of the data. In addition, it can be utilized for multimodal data fusion and geometric registration of shapes with hierarchical structure.

Acknowledgments and Disclosure of Funding

We thank the reviewers for their important comments and suggestions, and we thank Thomas Dagès for the helpful discussion. The work of YEL and RT was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 802735-ERC-DIFFOP. The work of YK was supported by the National Institutes of Health R01GM131642, UM1PA05141, P50CA121974, and U01DA053628.

References

- [1] B. Afsari. Riemannian l^p center of mass: existence, uniqueness, and convexity. *Proceedings of the American Mathematical Society*, 139(2):655–673, 2011.
- [2] D. Alvarez-Melis, Y. Mroueh, and T. Jaakkola. Unsupervised hierarchy matching with optimal transport over hyperbolic spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 1606–1617. PMLR, 2020.
- [3] Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *CR Acad. Sci. Paris Sér. I Math.*, 305:805–808, 1987.
- [4] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [5] R. Chakraborty, D. Seo, and B. C. Vemuri. An efficient exact-pga algorithm for constant curvature manifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3976–3984, 2016.
- [6] I. Chami, Z. Ying, C. Ré, and J. Leskovec. Hyperbolic graph convolutional neural networks. *Advances in neural information processing systems*, 32:4868–4879, 2019.
- [7] N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1853–1865, 2016.

- [8] C. Curtis, S. P. Shah, S.-F. Chin, G. Turashvili, O. M. Rueda, M. J. Dunning, D. Speed, A. G. Lynch, S. Samarajiwa, Y. Yuan, et al. The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. *Nature*, 486(7403):346–352, 2012.
- [9] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26:2292–2300, 2013.
- [10] J. Euzenat and P. Shvaiko. Ontology matching: State of the art and future challenges. *IEEE Transactions on knowledge and data engineering*, 25(1):158–176, 2013.
- [11] L. Falorsi, P. de Haan, T. R. Davidson, and P. Forré. Reparameterizing distributions on lie groups. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3244–3253. PMLR, 2019.
- [12] R. Flamary and N. Courty. Pot python optimal transport library. *GitHub: <https://github.com/rflamary/POT>*, page 144, 2017.
- [13] M. Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. In *Annales de l’institut Henri Poincaré*, volume 10, pages 215–310, 1948.
- [14] O. Ganea, G. Bécigneul, and T. Hofmann. Hyperbolic entailment cones for learning hierarchical embeddings. In *International Conference on Machine Learning*, pages 1646–1655. PMLR, 2018.
- [15] O. Ganea, G. Becigneul, and T. Hofmann. Hyperbolic neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, page 5350–5360, 2018.
- [16] I. Gohberg, P. Lancaster, and L. Rodman. Matrices and indefinite scalar products. 1983.
- [17] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
- [18] J. C. Gower, G. B. Dijkstra, et al. *Procrustes problems*, volume 30. Oxford University Press on Demand, 2004.
- [19] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- [20] D. Groisser. Newton’s method, zeroes of vector fields, and the riemannian center of mass. *Advances in Applied Mathematics*, 33(1):95–135, 2004.
- [21] D. N. Hayes, S. Monti, G. Parmigiani, C. B. Gilks, K. Naoki, A. Bhattacharjee, M. A. Socinski, C. Perou, and M. Meyerson. Gene expression profiling reveals reproducible human lung adenocarcinoma subtypes in multiple independent patient cohorts. *Journal of Clinical Oncology*, 24(31):5079–5090, 2006.
- [22] A. Hoyos-Idrobo. Aligning hyperbolic representations: an optimal transport-based approach. *arXiv preprint arXiv:2012.01089*, 2020.
- [23] J. J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- [24] H. Karcher. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics*, 30(5):509–541, 1977.
- [25] A. Klimovskaia, D. Lopez-Paz, L. Bottou, and M. Nickel. Poincaré maps for analyzing complex hierarchies in single-cell data. *Nature communications*, 11(1):1–9, 2020.
- [26] D. Koboldt, R. Fulton, M. McLellan, H. Schmidt, J. Kalicki-Veizer, J. McMichael, L. Fulton, D. Dooling, L. Ding, E. Mardis, et al. Comprehensive molecular portraits of human breast tumours. *Nature*, 490(7418):61–70, 2012.
- [27] D. Kohli, A. Cloninger, and G. Mishne. LDLE: Low distortion local eigenmaps. In *ICLR 2021 Workshop on Geometrical and Topological Representation Learning*, 2021.

- [28] I. Korsunsky, N. Millard, J. Fan, K. Slowikowski, F. Zhang, K. Wei, Y. Baglaenko, M. Brenner, P.-r. Loh, and S. Raychaudhuri. Fast, sensitive and accurate integration of single-cell data with harmony. *Nature methods*, 16(12):1289–1296, 2019.
- [29] A. Kumagai and T. Iwata. Unsupervised domain adaptation by matching distributions based on the maximum mean discrepancy via unilateral transformations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4106–4113, 2019.
- [30] M. Law, R. Liao, J. Snell, and R. Zemel. Lorentzian distance learning for hyperbolic representations. In *International Conference on Machine Learning*, pages 3672–3681. PMLR, 2019.
- [31] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [32] Q. Liu, M. Nickel, and D. Kiela. Hyperbolic graph neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 8228–8239, 2019.
- [33] A. Lou, I. Katsman, Q. Jiang, S. Belongie, S.-N. Lim, and C. De Sa. Differentiating through the fréchet mean. In *International Conference on Machine Learning (ICML)*, pages 6393–6403. PMLR, 2020.
- [34] B. Luo and E. Hancock. Feature matching with procrustes alignment and graph editing. 1999.
- [35] H. Maron, N. Dym, I. Kezurer, S. Kovalsky, and Y. Lipman. Point registration via efficient convex relaxation. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016.
- [36] R. J. McCann. Polar factorization of maps on riemannian manifolds. *Geometric & Functional Analysis GAFA*, 11(3):589–608, 2001.
- [37] H. H. Milioli, R. Vimeiro, I. Tishchenko, C. Riveros, R. Berretta, and P. Moscato. Iteratively refining breast cancer intrinsic subtypes in the metabric dataset. *BioData mining*, 9(1):1–8, 2016.
- [38] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l’Académie Royale des Sciences de Paris*, 1781.
- [39] Y. Nagano, S. Yamaguchi, Y. Fujita, and M. Koyama. A wrapped normal distribution on hyperbolic space for gradient-based learning. In *International Conference on Machine Learning (ICML)*, pages 4693–4702. PMLR, 2019.
- [40] I. Naim, Y. Song, Q. Liu, H. Kautz, J. Luo, and D. Gildea. Unsupervised alignment of natural language instructions with video segments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014.
- [41] M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems (NIPS)*, page 6341–6350, 2017.
- [42] M. Nickel and D. Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International Conference on Machine Learning (ICML)*, pages 3779–3788. PMLR, 2018.
- [43] V. Nygaard, E. A. Rødland, and E. Hovig. Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. *Biostatistics*, 17(1):29–39, 2016.
- [44] X. Pennec. Barycentric subspace analysis on manifolds. *The Annals of Statistics*, 46(6A):2711–2746, 2018.
- [45] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [46] J. G. Ratcliffe, S. Axler, and K. Ribet. *Foundations of hyperbolic manifolds*, volume 149. Springer, 1994.

- [47] P. L. C. Rodrigues, C. Jutten, and M. Congedo. Riemannian procrustes analysis: Transfer learning for brain–computer interfaces. *IEEE Transactions on Biomedical Engineering*, 66(8):2390–2401, 2018.
- [48] U. Shaham, K. P. Stanton, J. Zhao, H. Li, K. Raddassi, R. Montgomery, and Y. Kluger. Removal of batch effects using distribution-matching residual networks. *Bioinformatics*, 33(16):2539–2546, 2017.
- [49] D. Spohr, L. Hollink, and P. Cimiano. A machine learning approach to multilingual and cross-lingual ontology matching. In *International Semantic Web Conference*, pages 665–680. Springer, 2011.
- [50] T. Stuart, A. Butler, P. Hoffman, C. Hafemeister, E. Papalexi, W. M. Mauck III, Y. Hao, M. Stoeckius, P. Smibert, and R. Satija. Comprehensive integration of single-cell data. *Cell*, 177(7):1888–1902, 2019.
- [51] P. Tabaghi and I. Dokmanic. On procrustes analysis in hyperbolic space. *arXiv preprint arXiv:2102.03723*, 2021.
- [52] C. Wang and S. Mahadevan. Manifold alignment using procrustes analysis. In *International Conference on Machine Learning (ICML)*, pages 1120–1127, 2008.
- [53] O. Yair, M. Ben-Chen, and R. Talmon. Parallel transport on the cone manifold of spd matrices for domain adaptation. *IEEE Transactions on Signal Processing*, 67(7):1797–1811, 2019.
- [54] O. Yair, F. Dietrich, R. Talmon, and I. G. Kevrekidis. Domain adaptation with optimal transport on the manifold of spd matrices. *arXiv e-prints*, pages arXiv–1906, 2019.