

Horoball Hulls and Extents in Positive Definite Space^{*}

P. Thomas Fletcher, John Moeller, Jeff M. Phillips, and
Suresh Venkatasubramanian

University of Utah

Abstract. The space of positive definite matrices $P(n)$ is a Riemannian manifold with variable nonpositive curvature. It includes Euclidean space and hyperbolic space as submanifolds, and poses significant challenges for the design of algorithms for data analysis. In this paper, we develop foundational geometric structures and algorithms for analyzing collections of such matrices. A key technical contribution of this work is the use of *horoballs*, a natural generalization of halfspaces for non-positively curved Riemannian manifolds. We propose generalizations of the notion of a convex hull and a centerpoint and approximations of these structures using horoballs and based on novel decompositions of $P(n)$. This leads to an algorithm for approximate hulls using a generalization of extents.

1 Introduction

Data analysis and Euclidean geometry have traditionally been strongly linked, by representing data as points in a Euclidean space and comparing data using the Euclidean distance. However, as models for data analysis grow more sophisticated, it is becoming clearer that accurate modeling of data requires the use of non-Euclidean geometry and the induced geodesic distances. This geometry might be as simple as a surface embedded in a Euclidean space but in general may be represented as a Riemannian manifold with variable curvature [4].

One such manifold is $P(n)$, the manifold of real symmetric positive definite matrices. There are many application areas where the basic objects of interest, rather than points in Euclidean space, are elements of $P(n)$. In diffusion tensor imaging [3], matrices in $P(3)$ model the flow of water at each voxel of a brain scan, and a goal is to *cluster* these matrices into groups that capture common flow patterns along fiber tracts. In mechanical engineering [11], stress tensors are modeled as elements of $P(6)$, and identifying groups of similar tensors helps locate homogeneous regions in a material from samples. Kernel matrices in machine learning are elements of $P(n)$ [26], and motivated by the problems of learning and approximating kernels for machine learning tasks, there has been recent interest in studying the geometry of $P(n)$ and related spaces [20, 7, 28].

In all these areas, a problem of great interest is the analysis [15, 16] of collections of such matrices (finding central points, clustering, doing regression). It is important to note that in all these examples, the *Riemannian* structure of $P(n)$ is a crucial element of the modelling process: merely treating the matrices as points in \mathbb{R}^{n^2} and endowing the space with the Euclidean distance¹ does not capture the correct notion of distance

^{*} This research was supported in part by NSF awards SGER-0841185 and CCF-0953066 and a subaward to the University of Utah under NSF award 0937060 to CRA.

¹ In other words, computing the Frobenius distance between matrices.

or closeness that is meaningful in these applications [22, 21, 27, 25, 23]. For example, if we want to interpolate between two matrices of similar volume (determinant), a line between the two in \mathbb{R}^{n^2} contains matrices of volume well outside the range of volumes of the two matrices, but all matrices on a Riemannian geodesic have volume in the range.

Performing data analysis in such spaces requires the standard geometric toolkit that has proven successful for Euclidean data: methods for summarizing data sets (extents and core sets), finding representatives (centerpoints), and even performing accurate sampling (VC dimension estimates and ε -samples). In order to compute these objects, we need appropriate generalizations of the equivalent concepts in Euclidean spaces.

Our Contributions. In this paper we initiate a study of geometric algorithms on $P(n)$. We develop appropriate generalizations of halfspaces and convex hulls and also prove bounds on the VC-dimension of associated range spaces. We apply these results to the problem of estimating approximate extents for collections of matrices in $P(n)$, as well as studying approximate center points for such collections. Our results indicate that the *horoball* (a generalization of a halfspace) retains many, but not all, of the same combinatorial and structural properties of halfspaces, and is therefore a crucial building block for designing algorithms in this space.

1.1 Hulls and Convexity: From \mathbb{R}^d to $P(n)$

$P(n)$ is a good model space for algorithmic analysis. Like other Cartan-Hadamard (C-H) manifolds [8], its metric balls under the geodesic distance are (geodesically) convex, which is an important property². $P(n)$ has been extensively studied analytically, and there are simple closed form expressions for geodesics, the geodesic distance and other important constructs (see Section 2 for details). This is in contrast to other negatively-curved spaces like the δ -hyperbolic spaces [17], where in general one must assume some oracle to compute distances between points.

However, the variable curvature of $P(n)$ poses significant challenges for the construction of standard geometric primitives. A natural notion of a halfspace that is both “flat” and “convex” seems elusive, and it is not even clear whether a finite description of the geodesic convex hull of 3 points exists! To understand why this is the case, and why generalizing convex hulls to $P(n)$ is difficult, it is helpful to first understand the key properties used to define these notions in Euclidean space. In general, we desire a compact representation of the convex hull of a set of points. Call this property (C):

(C): Given a finite set of points X , the convex hull $\mathcal{C}(X)$ of X has a finite description using simple convex objects.

Euclidean hulls. In Euclidean space, the “simple convex objects” are halfspaces, which satisfy two properties:

(P1) The complement of a closed halfspace is an open halfspace; both are convex. Their boundary, a hyperplane, is also convex.

² Compare this with even simple positively curved manifolds like the n -dimensional sphere, in which this is no longer true.

(P2) Generically, d hyperplanes intersect at a single point and n hyperplanes partition \mathbb{R}^d into $\Theta(n^d)$ regions.

These two properties can be used to construct the convex hull (satisfying property (C)) efficiently in two different ways. *Painting Segments*: Given a finite set $X_0 \subset \mathbb{R}^d$ we can paint segments between all $x_1, x_2 \in X_0$; the union of these segments yields a set X_1 . We can recursively apply this procedure d times to generate $\langle X_2, \dots, X_d \rangle$, where X_d is the convex hull of X_0 [5]. This is the convex combination of X_0 , and the boundary is partitioned into a finite number of faces uniquely determined by painting segments among d -point subsets of X_0 .

Intersection of Convex Families: The convex hull of $X \subset \mathbb{R}^d$ is the intersection of all halfspaces which contain X ; we only need to consider the finite set of halfspaces supported by d points. We can also define the convex hull of X as the intersection of all balls which contain X . Again we only need to consider balls supported by d points; fixing this incidence, let their radius grow to infinity so they become halfspaces in the limit; see Figure 1(a).

Hyperbolic hulls. The first space we encounter as we move beyond Euclidean space towards $P(n)$ is the hyperbolic space \mathbb{H}^d , which is a Riemannian manifold of constant negative curvature. It is convenient to embed \mathbb{H}^d in the unit ball of \mathbb{R}^d using well-known models, specifically the Klein model [8, I.6] of \mathbb{H}^d in which geodesics are straight lines, and the Poincaré model [8, I.6] of \mathbb{H}^d in which metric balls are Euclidean balls and geodesics are circular arcs normal to the boundary of the unit ball.

Since in the Klein model, geodesics are straight lines, the *painting segments* construction yields a convex hull in the same way as in Euclidean space; see Figure 1(b). Similarly, a *hyperbolic halfspace* can be written as the intersection of the unit ball with a Euclidean halfspace, and so a finite description of the convex hull can be obtained via the intersection of halfspaces; see Figure 1(c).

Hulls in $P(n)$. Once we reach $P(n)$, these concepts break down. There is no way, in general, to construct a halfspace (convex, and whose complement is convex) supported by d points; such an object is called a *totally geodesic submanifold* and might not pass through any given set of d points. This rules out constructions via the intersection of convex families. Constructing a hull via painting of segments also does not work; the




Fig. 1. Illustration of different constructions of hulls in \mathbb{R}^2 (a) and \mathbb{H}^2 under the Klein model (b,c) and Poincaré model (d). (a) Intersection of halfspaces, one halfspace as limit of ball supported by 2 points. (b) Painting segments, X_0 as circles, X_1 as segments, and X_2 shaded. (c) Intersection of hyperbolic-halfspaces, forming the convex hull. (d) Intersection of horoballs, forming the ball hull.

resulting process may not terminate in a finite number of steps, and the resulting object might be full-dimensional, and would not in general lend itself to a finite description.

There is however another way to approach the idea of halfspaces. Returning to \mathbb{H}^d and the Poincaré model, consider a ball fixed at a point whose radius is allowed to grow to infinity. Such a ball is called a *horoball*, and is convex. In Euclidean space, this construction yields a halfspace passing through the fixed point, but in a curved space, the ball never completely “flattens” out in the sense of property (P1). Horoballs can be described finitely by a point and a tangent vector (in the same way a Euclidean halfspace can be described by a point and a normal). We can then describe the intersection of all horoballs containing X , which we call the *ball hull*; it is convex and contains the convex hull of X ; see Figure 1(d). In the rest of this paper, we will focus our attention on horoballs and the ball hull.

1.2 Technical Overview

A key conceptual insight in this work is that the horoball acts functionally like a halfspace, and can be used as a replacement for halfspace in spaces (like $P(n)$) which do not in general admit halfspaces that span arbitrary sets of points. As justification for this insight, our main result is an algorithm for computing an approximate ball hull of a set of points in $P(n)$, where the approximation uses a generalized notion of *extent* defined analogously to how (hyperplane) extent is defined in \mathbb{R}^d . The construction itself follows the rough outline of approximate extent constructions in \mathbb{R}^d . In fact, we exploit the fact that $P(n)$ admits a decomposition into a collection of Euclidean subspaces, each “indexed” by an element of $SO(n)$, the group of $n \times n$ rotation matrices, and use a grid construction to cover $SO(n)$ with a net, followed by building convex hulls in each of the (finitely many) Euclidean subspaces induced by the net and combining them. We expect that this decomposition will be of independent interest.

We announce two other uses of horoballs, with details in the full version. We can define range spaces of horoballs, and we analyze shatter dimension and VC-dimension of $P(2)$. We also use these results to study center points in $P(n)$.

1.3 Related Work

The mathematics of Riemannian manifolds, Cartan-Hadamard manifolds, and $P(n)$ is well-understood; the book by Bridson and Haefliger [8] is an invaluable reference on metric spaces of nonpositive curvature, and Bhatia [5] provides a detailed study of $P(n)$ in particular. However, there are very few *algorithmic* results for problems in these spaces. To the best of our knowledge, the only prior work on algorithms for positive definite space are the work by Moakher [22] on mean shapes in positive definite space, the work by Fletcher and Joshi [15] on principal geodesic analysis in symmetric spaces, the robust median algorithms of Fletcher *et al* [16] for general manifolds (including $P(n)$ and $SO(n)$), and the generic approximation technique of Arnaudon and Nielsen [2] for the Riemannian 1-center.

Geometric algorithms in hyperbolic space are much more tractable. The Poincaré and Klein models of hyperbolic space preserve different properties of Euclidean space, and many algorithm carry over directly with no modifications. Leibon and Letscher [19] were the first to study basic geometric primitives in general Riemannian manifolds,

constructing Voronoi diagrams and Delaunay triangulations for sufficiently dense point sets in these spaces. Their work was subsequently improved by Dyer *et al.* [13]. Eppstein [14] described hierarchical clustering algorithms in hyperbolic space.

δ -hyperbolic spaces [17] (metric spaces that “look” negatively curved without necessarily having a smooth notion of curvature) have also been studied. Krauthgamer and Lee [18] studied the nearest neighbor problem for points in δ -hyperbolic space. Chepoi *et al* [9, 10] advanced this line of research, providing algorithms for computing the diameter and minimum enclosing ball of collections of points in δ -hyperbolic space. Work by Billera *et al.* [6] showed how to model the space of phylogenetic trees as a specific CAT(0) space [8, II.1]; work by Owen and Provan investigated how to efficiently compute geodesics in such a space [24].

2 Preliminaries

In this section we present the basic algebra and geometry needed to understand our technical results without additional outside references. $P(n)$ is the manifold consisting of symmetric positive-definite real matrices. As a manifold, it has a Euclidean tangent space $T_pP(n)$ at each point p , represented by the space of symmetric matrices $S(n)$. A point in the tangent space represents a vector tangent to a curve that passes through the point p . The velocity of a particle moving along such a curve, for example, would be represented by a vector in $T_pP(n)$ when it passes through p .

A *geodesic* through a point p is a special curve determined entirely by a tangent vector $A \in T_pP(n)$. This relationship between the tangent space and the manifold is realized by the \exp map $\exp_p : S(n) \rightarrow P(n)$, defined as $\exp_p(A) = pe^{p^{-1}A}$, where e^X is just the matrix exponential. If $c(t)$ is a geodesic with tangent A at $c(0) = p$, then $c(t) = \exp_p(tA)$. For simplicity, we often assume that $p = I$ so $\exp_I(A) = e^A$. The \exp map has a simple intuition: in Euclidean space, $\exp_p(\mathbf{u}) = p + \mathbf{u}$; that is, we just move from the point p to another point that is in the direction of \mathbf{u} , $\|\mathbf{u}\|$ units away.

We can then measure distance between two points on the manifold by finding a geodesic between them, solving for the unknown tangent vector, and measuring its length. The \exp map is invertible, and its inverse is the \log map, $\log_p : P(n) \rightarrow S(n)$, given by $\log_p(q) = p \log(p^{-1}q)$, where \log is the inverse of the matrix exponential. $P(n)$ is also endowed with the special property that \exp_p is invertible across the entire manifold, letting us measure distance between any two points. Because $P(n)$ is a *Riemannian* manifold, $T_pP(n)$ has an inner product $\langle A, B \rangle_p = \text{tr}(p^{-1}Ap^{-1}B)$, and $\|A\|_p = \sqrt{\langle A, A \rangle_p}$. The resulting metric is then $D(p, q) = \|\log_p(q)\| = \sqrt{\text{tr}(\log(p^{-1}q)^2)}$. Usually we assume that $\|A\|_p = 1$, meaning that a geodesic’s parameter t equals distance traveled.

As mentioned before, $P(n)$ is a Riemannian manifold of non-positive curvature. Its \exp map is *surjective*, which enables us to talk about geodesics between any two points on the manifold. Therefore we need not concern ourselves with a radius of convexity; a convex subset of $P(n)$ need not be bounded. Another important property of $P(n)$ is that it is *symmetric*. This means that there is always an isometry that moves a point p to a point q without altering the metric properties of the manifold, offering the equivalent of translation invariance. A good reference for the reader is [8, II.10].

Structure of $P(2)$. Worth mentioning is the specific 3-dimensional manifold $P(2)$, which has a structure that we exploit frequently. $P(2)$ is isometric to $\mathbb{R} \times P(2)_1$, where

$P(2)_1$ is the same as \mathbb{H}^2 with a factor of $1/\sqrt{2}$ on the metric. $P(2)_1$ is so named because it represents the submanifold of $P(2)$ containing all p.d. matrices of determinant 1. Intuitively, we can think of this as a “stack” of hyperbolic spaces, leading us further to a cylindrical representation of the space (see Figure 2). Decomposing a matrix p as $(r, p') = e^{r/2} \cdot p'$, we can represent the (log) determinant of p as r . Using a polar representation, we can break p' down further and realize the *anisotropy*, or ratio of eigenvalues, as a radial coordinate. Since the eigenvectors form a rotation matrix, we take the angle of this rotation (times 2) as the remaining coordinate.

2.1 Busemann Functions and Horoballs

In \mathbb{R}^d , the convex hull of a finite set can be described by a finite number of hyperplanes each supported by d points from the set. A hyperplane through a point may also be thought of as the limiting case of a sphere whose center has been moved away to infinity while a point at its surface remains fixed. This notion of “pulling away to infinity” can be formalized: given a geodesic ray $c(t) : \mathbb{R}^+ \rightarrow M$ on a Cartan-Hadamard manifold M , a *Busemann function* $b_c : M \rightarrow \mathbb{R}$ is defined $b_c : p \mapsto \lim_{t \rightarrow \infty} D(p, c(t)) - t$. An important property of b_c is that it is convex [8, II.8].

As an example, we can easily compute the Busemann function in \mathbb{R}^n for a ray $c(t) = t\mathbf{u}$, where \mathbf{u} is a unit vector. Since $\lim_{t \rightarrow \infty} \frac{1}{2t} (\|p - t\mathbf{u}\| + t) = 1$,

$$b_c(p) = \lim_{t \rightarrow \infty} \frac{1}{2t} (\|p - t\mathbf{u}\|^2 - t^2) = -\langle p, \mathbf{u} \rangle.$$

A *horoball* $B_r(b_c) \subset M$ is a sublevel set of b_c ; that is, $B_r(b_c)$ is the set of all $p \in M$ such that $b_c(p) \leq r$ (recall Figure 1). Since b_c is convex, any sublevel set of it is convex, and hence any horoball is convex. Continuing with the example of Euclidean space, horoballs are simply halfspaces: all $p \in \mathbb{R}^d$ such that $-\langle p, \mathbf{u} \rangle \leq r$.

3 Ball Hulls

We now introduce our variant of the convex hull in $P(n)$, which we call the ball hull. For a subset $X \subset P(n)$, the *ball hull* $\mathcal{B}(X)$ is the intersection of all horoballs that also contain X :

$$\mathcal{B}(X) = \bigcap_{b_c, r} B_r(b_c), \quad X \subset B_r(b_c).$$

Properties of the ball hull. Recall that the ball hull can be seen as an alternate generalization of the Euclidean convex hull (i.e. via intersection of halfspaces) to $P(n)$. Furthermore, since it is the intersection of closed convex sets, it is itself guaranteed to be closed and convex (and therefore $\mathcal{C}(X) \subseteq \mathcal{B}(X)$). We can also show that it shares critical parts of its boundary with the convex hull (Theorem 1), but unfortunately, we cannot represent it as a finite intersection of horoballs (Theorem 2). We state these theorems here, and defer proofs to the full version.


Theorem 1. *Every $x \in X$ (X finite) on the boundary of $\mathcal{B}(X)$ is also on the boundary of $\mathcal{C}(X)$ (i.e., $X \cap \partial \mathcal{B}(X) \subseteq X \cap \partial \mathcal{C}(X)$).*

Theorem 2. *In general, the ball hull cannot be described as the intersection of a finite set of horoballs.*

3.1 The ε -Ball Hull

Theorem 2 indicates that we cannot maintain a finite representation of a ball hull. However, as we show in this section, we can maintain a finite-sized *approximation* to the ball hull. Our approximation will be in terms of *extents*; intuitively, a set of horoballs approximates the ball hull if a geodesic traveling in any direction traverses approximately the same distance inside the ball hull as it does inside the approximate hull.

In Euclidean space, we can capture extent by measuring the distance between two parallel hyperplanes that sandwich the set. The analogue to this construction in $P(n)$ is the *horoextent*, the distance along a geodesic between two opposing horoballs. Let $c(t) = qe^{tq^{-1}A}$ be a geodesic, and $X \subset P(n)$. The *horoextent* $E_c(X)$ with respect to c is defined as:



horoextent

$$E_c(X) = \left| \max_{p \in X} b_{c^+}(p) + \max_{p \in X} b_{c^-}(p) \right|,$$

where b_{c^+} is the Busemann function created when we follow $c^+(t) = c(t)$ to infinity as normal, while b_{c^-} is the Busemann function created when we follow the geodesic pointing in the opposite direction, $c^-(t) = qe^{tq^{-1}(-A)} = c(-t)$. Stated differently: $b_{c^+}(p) = \lim_{t \rightarrow +\infty} (D(c(t), p) - t)$, and $b_{c^-}(p) = \lim_{t \rightarrow -\infty} (D(c(t), p) + t)$. Observe that for any c , $E_c(X) = E_c(\mathcal{C}(X)) = E_c(\mathcal{B}(X))$.

In Euclidean space, the distance between parallel planes is a constant. In general, because of the effects of curvature, the distance between horoballs depends on the geodesic used. For instance in $P(n)$, horofunctions are nonlinear, so the distance between opposing horoballs is not constant. The width of the intersection of the opposing horoballs is taken along the geodesic c , and a geodesic is described by a point q and a direction A . We fix the point q so that we need only choose a uniform grid of directions A for our approximation.

An intersection of horoballs is called an ε -ball hull with origin q ($\mathcal{B}_{\varepsilon,q}(X)$) if for all geodesic rays c such that $c(0) = q$, $|E_c(\mathcal{B}_{\varepsilon,q}(X)) - E_c(X)| \leq \varepsilon$. For convenience, we assume that $I \in \mathcal{C}(X)$ (this assumption will be removed in the full version), and our origin $q = I$. Then we will refer to $\mathcal{B}_{\varepsilon,I}(X)$ as just an ε -ball hull $\mathcal{B}_\varepsilon(X)$.

We will use $D_X \leq \text{diam}(X) = \max_{p,q \in X} D(p,q)$ in our bounds; see the full version for more precise definition, and note $\text{diam}(X)$ is an intrinsic parameter of the data.

4 Constructing the ε -Ball Hull

Main result. In this section we construct a finite-sized ε -ball hull.

Theorem 3. Let $\Gamma_n(\varepsilon, D_X) = (\sinh(D_X)/\varepsilon)^{n-1}$. For a set $X \subset P(n)$ of size N (for constant n), we can construct an ε -ball hull of size $O(\Gamma_n(\varepsilon, D_X) \cdot (D_X/\varepsilon)^{(n-1)/2})$ in time $O(\Gamma_n(\varepsilon, D_X) \cdot ((D_X/\varepsilon)^{n-3/2} + N))$. Furthermore, we can construct a coreset $Y \subset X$ of size $O(\Gamma_n(\varepsilon, D_X) \cdot (D_X/\varepsilon)^{(n-1)/2})$ whose $(\varepsilon/2)$ -ball hull is an ε -ball hull of X .

Proof overview. We make extensive use of a structural decomposition in our proof. But first it will be helpful to define a *flat*. Let us define a *subspace* of a manifold as the result of applying \exp_p to each point of a subspace of the tangent space $T_pP(n)$. If the resulting submanifold is isometric to a Euclidean space, then it is called a *flat*. A flat has




Fig. 2. (a) $P(2)$ as revolving door, (b) projection of $X \subset P(2)$ onto $\det(x) = 1$, (c) $X \subset P(2)$, (d) a flat in $P(2)$, (e) flat of $P(2)$ under $\log(\cdot)$ map. Two horospheres drawn in views (b-d).

the important property of being convex (in general a subspace is not). One canonical example of a flat is the subspace of positive sorted diagonal matrices.

$P(n)$ can then be realized as the union of a set of n -dimensional flats, and the space can be parameterized by a rotation matrix $Q \in SO(n)$, one for each flat F . The intersection of these flats is the line of multiples of I . In $P(2)$, we can picture these flats as panes in a revolving door; see Figure 2(a).

We construct ε -ball hulls by discretizing $P(n)$ in two steps. First, we show that within a flat F (i.e., given a rotation $Q \in SO(n)$) we can find a finite set of minimal horoballs exactly, or we can use ε -kernel machinery [1] to approximate this structure. This is done by showing an equivalence between halfspaces in F and horoballs in $P(n)$ in Section 4.1. This result implies that computing all minimal horoballs with respect to a rotation Q is equivalent to computing a convex hull in Euclidean space.

Second, we show that instead of searching over the entire space of rotations $SO(n)$, we can discretize it into a finite set of rotations such that when we calculate the horoballs with respect to each of these rotations, the horoextents of the resulting ε -ball hull are not too far from those of the ball hull. In order to do this, we prove a Lipschitz bound for horofunctions (and hence horoextents) on the space of rotations.

Proving this theorem is quite technical. We first prove a Lipschitz bound in $P(2)$, where the space of rotations is a circle (as in Figure 2(b)). After providing a bound in $P(2)$ we decompose the distance between two rotations in $SO(n)$ into $\lfloor n/2 \rfloor$ angles defined by 2×2 submatrices in an $n \times n$ matrix. In this setting it is possible to apply the $P(2)$ Lipschitz bound $\lfloor n/2 \rfloor$ times to get the full bound. We present the proof for $P(2)$ in Section 4.2, and the generalization to $P(n)$ in Section 4.3. Finally, we combine these results in an algorithm in Section 4.4.

4.1 Decomposing $P(n)$ into Flats

A critical operation associated with the decomposition (illustrated in Figure 2) is the *horospherical projection function* $\pi_F : P(n) \rightarrow F$ that maps a point $p \in P(n)$ to a point $\pi_F(p)$ in a n -dimensional flat F . For each Busemann function b_c there exists a flat F for which it is invariant under the associated projection π_F ; that is, $b_c(p) = b_c(\pi_F(p))$ for all $p \in P(n)$. Using π_F for associated geodesic $c(t) = e^{tA}$ where $A \in S(n)$, the Busemann function $b_c : P(n) \rightarrow \mathbb{R}$ can be written [8, II.10]

$$b_c(p) = -\text{tr}(A \log(\pi_F(p))).$$

It is irrelevant which point is chosen for the origin of the geodesic ray, so we usually assume that it is chosen in such a way that $b_c(I) = 0$ in $P(n)$.

In $P(2)$ it is convenient to visualize Busemann functions through horospheres. We can embed $P(2)$ in \mathbb{R}^3 where the log of the determinant of elements grows along one

axis. The orthogonal planes contain a model of hyperbolic space called the *Poincaré disk* that is modeled as a unit disk, with boundary at infinity represented by the unit circle. Thus the entire space can be seen as a cylinder, as shown in Figure 2(c). Within each cross section with constant determinant (Figure 2(b)), the horoballs are disks tangent to the boundary at infinity. Within each flat F (Figure 2(d)) under a $\log(\cdot)$ map (Figure 2(e)) the horoballs are halfspaces. The full version provides a more technical treatment of this decomposition and proofs of technical lemmas.

Rotation of Busemann functions. The following lemma describes how geodesics (and horofunctions) are transformed by a rotation. In particular, this allows us to pick a flat where computation of b_c is convenient, and rotate the point set by Q to compute b_c instead of attempting computation of $b_{c'}$ directly.

Lemma 1. *For $p \in P(n)$, rotation matrix Q , geodesics $c(t) = e^{tA}$ and $c'(t) = e^{tQAQ^T}$, then $b_{c'}(p) = b_c(Q^T p Q)$.*

Projection to k -flat. We now establish an equivalence between horoballs and halfspaces. That is, after we compute the projection of our point set, we can say that the point set X lies inside a horoball $B_r(b_c)$ if and only if its projection $\pi_F(X)$ lies inside a halfspace H_r of F (recall that F is isometric to a Euclidean space under \log).

Lemma 2. *For any horoball $B_r(b_c)$, there is a halfspace $H_r \subset \log(F) \subset S(n)$ such that $\log(\pi_F(B_r(b_c))) = H_r$.*

Proof. If $b_c(p) \leq r$, $p \in P(n)$, and $c(t) = e^{tA}$, then $-\text{tr}(A \log(\pi_F(p))) \leq r$. Since $\pi_F(p)$ is positive-definite, $\log(\pi_F(p))$ is symmetric. But $\text{tr}((\cdot)(\cdot))$ defines an inner product on the Euclidean space of symmetric $n \times n$ matrices. Then the set of all Y such that $-\text{tr}(AY) \leq r$ defines a halfspace H_r whose boundary is perpendicular to A . Furthermore, given a matrix v such that $b(vp v^T) \leq r$ (see the full version), we can compute $v e^Y v^T$ for every $Y \in H_r$, so every such Y maps back to an element of $B_r(b_c)$. \square

4.2 A Lipschitz bound in $P(2)$

To show our Lipschitz bound we analyze the deviation between two flats with similar directions. Since any two flats F and F' are identified with rotations Q and Q' , we can move a point from F to F' simply by applying the rotation $Q^T Q'$, and measure the angle θ between the flats. If we consider a geodesic $c \subset F$ we can apply $Q^T Q'$ to c to get c' , then for any point $p \in P(n)$ we bound $|b_c(p) - b_{c'}(p)|$ as a function of θ . Technical proofs are in the full version.

Rotations in $P(2)$. We start with some technical lemmas that describe the locus of rotating points in $P(2)$.

Lemma 3. *Given a rotation matrix $Q \in SO(2)$ corresponding to an angle of $\theta/2$, Q acts on a point $p \in P(2)$ via QpQ^T as a rotation by θ about the (geodesic) axis $e^I = e^I$.*

By Lemma 3, as we apply a rotation to p , it moves in a circle. Because any rotation Q has determinant 1, $\det(QpQ^T) = \det(p)$. This leads to the following corollary:

Corollary 1. *In $P(2)$, the radius of the circle that p travels on is $D(\sqrt{\det(p)}I, p)$. Such a circle lies entirely within a submanifold of constant determinant.*

In fact, any submanifold $P(2)_r$ of points with determinant equal to some $r \in \mathbb{R}^+$ is isometric to any other such submanifold $P(2)_s$ for $s \in \mathbb{R}^+$. This is easily seen by considering the distance function $\text{tr}(\log(p^{-1}q))$; the determinants of p and q will cancel. We pick a natural representative of these submanifolds, $P(2)_1$. This submanifold forms a complete metric space of its own that has special structure:

Lemma 4. $P(2)_1$ has constant sectional curvature $-1/2$.

To bound the error incurred by discretizing the space of directions, we need to understand the behavior of b_c as a function of a rotation Q . We show that the derivative of a geodesic is constant on $P(n)$.

Lemma 5. For a geodesic ray $c(t) = e^{tA}$ where $\|A\| = 1$, then $\|\nabla b_c\| = 1$ at any point $p \in P(n)$.

Lipschitz condition on Busemann functions in $P(2)$.

Theorem 4. Consider $Q \in SO(2)$ corresponding to $\theta/2$, $c(t) = e^{tA}$, and $c'(t) = e^{tQ^T A Q}$. Then for any $p \in X$

$$|b_c(p) - b_{c'}(p)| \leq |\theta| \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right).$$

Proof. The derivative of a function f along a curve $\gamma(t)$ has the form $\langle \nabla f|_{\gamma(t)}, \gamma'(t) \rangle$, and has greatest magnitude when the tangent vector $\gamma'(t)$ to the curve and the gradient $\nabla f|_{\gamma(t)}$ are parallel. When this happens, the derivative reaches its maximum at $\|\nabla f|_{\gamma(t)}\| \cdot \|\gamma'(t)\|$. Since $\|\nabla b_c\| = 1$ anywhere by Lemma 5, the derivative of b_c along γ at $\gamma(t)$ is bounded by $\|\gamma'(t)\|$. We are interested in the case where $\gamma(\theta)$ is the circle in $P(2)$ defined by tracing $Q(\theta/2)pQ(\theta/2)^T$ for all $-\pi < \theta \leq \pi$. By Corollary 1, we know that this circle has radius $D(\sqrt{\det(p)}I, p) \leq D(I, p) \leq D_X$ and lies entirely within a submanifold of constant determinant, which by Lemma 4 also has constant curvature $\kappa = -1/2$. This implies that

$$\|\gamma'(\theta)\| = \frac{1}{\sqrt{-\kappa}} \sinh(\sqrt{-\kappa}r) = \sqrt{2} \sinh\left(\frac{D(\sqrt{\det(p)}I, p)}{\sqrt{2}}\right) \leq \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right)$$

for any value of $\theta \in (-\pi, \pi]$ [8, I.6]. Then

$$|b_c(p) - b_{c'}(p)| = |b_c(p) - b_c(Q^T p Q)| \leq |\theta| \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right). \quad \square$$

4.3 Generalizing to $P(n)$

Now to generalize to $P(n)$ we need to decompose the projection operation $\pi_F(\cdot)$ and the rotation matrix Q . We can compute π_F recursively, and it turns out that this fact helps us to break down the analysis of rotations. Since we can decompose any rotation into a product of $\lfloor n/2 \rfloor$ 2×2 rotation matrices, decomposing the computation of π_F in a similar manner lets us build a Lipschitz condition for $P(n)$. The full version formalizes rotation decompositions we use.

Theorem 5 (Lipschitz condition on Busemann functions in $P(n)$). Consider a set $X \subset P(n)$, a rotation matrix $Q \in SO(n)$ corresponding to an angle $\theta/2$, geodesics $c(t) = e^{tA}$ and $c'(t) = e^{tQ A Q^T}$. Then for any $p \in X$

$$|b_c(p) - b_{c'}(p)| \leq |\theta| \cdot \left\lfloor \frac{n}{2} \right\rfloor \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right).$$

Proof. Every rotation Q may be decomposed into a product of rotations, relative to some orthonormal basis B ; that is, $Q = B(Q_1 Q_2 \cdots Q_{k-1} Q_k) B^T$ where $k = \lfloor n/2 \rfloor$ and Q_i is a 2×2 subblock rotation corresponding to an angle $\theta_i/2$ with $|\theta_i| \leq |\theta|$. Now applying Lemma 1, we can factor out B : $c'(t) = e^{tQAQ^T} = e^{tB(Q_1 \cdots Q_k) B^T A B (Q_k^T \cdots Q_1^T) B^T}$, and let $\hat{c}'(t) = e^{t(Q_1 \cdots Q_k) \hat{A} (Q_k^T \cdots Q_1^T)}$, where $\hat{A} = B^T A B$. This means that $b_{c'}(p) = b_{c'}(B^T p B) = b_{\hat{c}'}(\hat{p})$ for $\hat{p} = B^T p B$. Also, since $c(t) = e^{tA} = e^{tB \hat{A} B^T}$, $b_c(p) = b_{\hat{c}}(\hat{p})$ for $\hat{c}(t) = e^{t \hat{A}}$. From this point we will omit the “hat” notation and just assume the change of basis.

$$\text{Then } |b_c(p) - b_{c'}(p)| = \left| \sum_{i=1}^k (b_{c'}^{i-1}(p) - b_{c'}^i(p)) \right| \leq \sum_{i=1}^k |b_{c'}^{i-1}(p) - b_{c'}^i(p)|,$$

where $b_{c'}^0(p) = b_c(p)$ and $b_{c'}^i(p)$ is $b_c(p)$ with the first i rotations successively applied, so $b_{c'}^i(p) = b_c(Q_i^T \cdots Q_1^T p Q_1 \cdots Q_i)$. Then by Theorem 4 $|b_{c'}^{i-1}(p) - b_{c'}^i(p)| \leq |\theta_i| \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right)$, and therefore, since for all i we have $|\theta_i| \leq |\theta|$,

$$|b_c(p) - b_{c'}(p)| \leq \left(\sum_{i=1}^k |\theta_i| \right) \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right) \leq |\theta| \cdot \left\lfloor \frac{n}{2} \right\rfloor \cdot \sqrt{2} \sinh\left(\frac{D_X}{\sqrt{2}}\right). \quad \square$$

4.4 Algorithm

For $X \subset \mathbb{P}(n)$ we can construct ε -ball hull as follows. We place a grid G_ε on $\text{SO}(n)$ so that for any $Q' \in \text{SO}(n)$, there is another $Q \in G_\varepsilon$ such that the angle between Q and Q' is at most $(\varepsilon/4)/(2\lfloor n/2 \rfloor \sqrt{2} \sinh(D_X/\sqrt{2}))$. For each $Q \in G_\varepsilon$, we consider $\pi_F(X)$, the projection of X into the associated n -flat F associated with Q . Within F , we build an $(\varepsilon/4D_X)$ -kernel [1] K_F of $\log(\pi_F(X))$, and return the horoball associated with the hyperplane passing through each facet of $\mathcal{C}(K_F)$ in F , using the transformation specified in Lemma 2. (This step can be replaced with an exact convex hull of $\log(\pi_F(X))$.) This produces a coreset of X for extents, $K = \bigcup_F K_F$ for all F associated with a $Q \in G_\varepsilon$.

To analyze this algorithm we can now consider any direction $Q' \in \text{SO}(n)$ and a horofunction $b_{c'}$ that lies in the associated flat F' . There must be another direction $Q \in G_\varepsilon$ such that the angle between Q and Q' is at most $(\varepsilon/4)/(2\lfloor n/2 \rfloor \sqrt{2} \sinh(D_X/\sqrt{2}))$. Let b_c be the similar horofunction to $b_{c'}$, except it lies in the flat F associated with Q . This ensures that for any point $p \in X$, we have $|b_{c'}(p) - b_c(p)| \leq \varepsilon/4$. Also, for $p \in X$, there is a $q \in K_F$ such that $b_c(p) - b_c(q) \leq (\varepsilon/4D_X) \cdot D_X = \varepsilon/4$. Thus $b_{c'}(p) - b_c(q) \leq \varepsilon/2$. Since $E_{c'}(X)$ depends on two points in X , and each point p changes at most $\varepsilon/2$ from $b_{c'}(p)$ to $b_c(q)$, we can argue that $|E_{c'}(X) - E_c(K)| \leq \varepsilon$. Since this holds for any direction $Q' \in \text{SO}(n)$, the returned set of horoballs defines an ε -ball hull.

Let $\Gamma_n(\varepsilon, D_X) = (\sinh(D_X)/\varepsilon)^{n-1}$. For constant n , the grid G_ε is size $O(\Gamma_n(\varepsilon, D_X))$. In each flat, the $(\varepsilon/4D_X)$ -kernel is designed to have convex hull with $O((\varepsilon/D_X)^{(n-1)/2})$ facets [1, 12], and can be computed in $O(N + (D_X/\varepsilon)^{n-3/2})$ time. Thus a ε -ball hull represented as the intersection of $O(\Gamma_n(\varepsilon, D_X) \cdot (D_X/\varepsilon)^{(n-1)/2})$ horoballs can be computed in $O(\Gamma_n(\varepsilon, D_X) \cdot ((D_X/\varepsilon)^{n-3/2} + N))$ time, proving Theorem 3.

Acknowledgements. We sincerely thank the anonymous reviewers for their many insightful comments and for their time digesting our work.

References

1. P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *JACM*, 51, 2004.

2. M. Arnaudon and F. Nielsen. On approximating the Riemannian 1-center. *arXiv:1101.4718v1*, 2011.
3. P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophys. J.*, 66(1):259–267, 1994.
4. M. Belkin and P. Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56(1):209–239, 2004.
5. R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2006.
6. L. Billera, S. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4):733–767, 2001.
7. S. Bonnabel and R. Sepulchre. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM J. on Matr. Anal. & App.*, 31(3):1055–1070, 2010.
8. M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer, 2009.
9. V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. Diameters, centers, and approximating trees of delta-hyperbolic geodesic spaces and graphs. *SoCG*, 2008.
10. V. Chepoi and B. Estellon. Packing and covering δ -hyperbolic spaces by balls. *LNCS; Vol. 4627*, 2007.
11. S. Cowin. The structure of the linear anisotropic elastic symmetries. *Journal of the Mechanics and Physics of Solids*, 40:1459–1471, 1992.
12. R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Journal of Approximation Theory*, 10:227–236, 1974.
13. R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. In *SGP*, pages 1393–1402, 2008.
14. D. Eppstein. Squarepants in a tree: Sum of subtree clustering and hyperbolic pants decomposition. *ACM Transactions on Algorithms (TALG)*, 5, 2009.
15. P. T. Fletcher and S. Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *Com Vis & Math Methods Med Biomed Im Anal*, pages 87–98, 2004.
16. P. T. Fletcher, S. Venkatasubramanian, and S. Joshi. The geometric median on Riemannian manifolds with application to robust atlas estimation. *NeuroImage*, 45:S143–52, 2009.
17. M. Gromov. Hyperbolic groups. *Essays in group theory*, 8:75–263, 1987.
18. R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. *FOCS*, 2006.
19. G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, 2000.
20. G. Meyer, S. Bonnabel, and R. Sepulchre. Regression on fixed-rank positive semidefinite matrices: a Riemannian approach. *arXiv:1006.1288*, 2010.
21. M. Moakher. A Differential Geometric Approach to the Geometric Mean of Symmetric Positive-Definite Matrices. *SIAM J. on Matr. Anal. & App.*, 26, 2005.
22. M. Moakher and P. G. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. *Visualization and Processing of Tensor Fields*, 2006.
23. Y. Nesterov and M. Todd. On the Riemannian geometry defined by self-concordant barriers and interior-point methods. *FoCM*, 2(4):333–361, 2008.
24. M. Owen and J. Provan. A Fast Algorithm for Computing Geodesic Distances in Tree Space. *arXiv:0907.3942*, 2009.
25. X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *IJCV*, 66(1):41–66, 2006.
26. J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge, 2004.
27. S. Smith. Covariance, subspace, and intrinsic Cramér-Rao bounds. *Signal Processing, IEEE Transactions on*, 53(5):1610–1630, 2005.
28. B. Vandereycken, P. Absil, and S. Vandewalle. A Riemannian geometry with complete geodesics for the set of positive semidefinite matrices of fixed rank. *status: submitted*, 2010.