Statistics of Shape via Principal Geodesic Analysis on Lie Groups

P. Thomas Fletcher, Conglin Lu and Sarang Joshi
Medical Image and Display Analysis Group,
University of North Carolina
fletcher@cs.unc.edu

Abstract

Principal component analysis has proven to be useful for understanding geometric variability in populations of parameterized objects. The statistical framework is well understood when the parameters of the objects are elements of a Euclidean vector space. This is certainly the case when the objects are described via landmarks or as a dense collection of boundary points. We have been developing representations of geometry based on the medial axis description or m-rep. Although this description has proven to be effective, the medial parameters are not naturally elements of a Euclidean space. In this paper we show that medial descriptions are in fact elements of a Lie group. We develop methodology based on Lie groups for the statistical analysis of medially-defined anatomical objects.

1. Introduction

Shape analysis is emerging as an important area of image processing and computer vision. Model-based approaches [4, 3, 21, 13] are popular due to their ability to robustly represent objects found in images. Principal component analysis (PCA) [10] is a prevalent technique for describing model variability. However, PCA is only applicable when model parameters are elements of a Euclidean vector space.

The focus of our research has been the application of shape analysis for medical image processing to improve both the accuracy of medical diagnosis as well as the understanding of processes behind growth and disease [5]. In our previous work [11] we have developed methodology based on medial descriptions called m-reps to quantify shape variability and explain it in intuitive terms such as local thickness, bending and widening.

In this paper we show that m-rep models are elements of a Lie group. We develop a framework that extends the notion of PCA to Lie groups, and we apply it to the statistical analysis of shape using medial representations. As the medial representation is fundamental to our analysis, we describe it briefly.

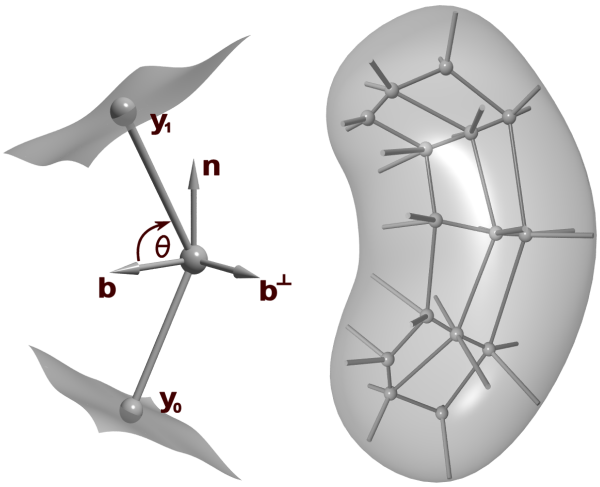

Figure 1. Medial atom with a cross-section of the boundary surface it implies (left). An m-rep model of a kidney and its boundary surface (right).

1.1. M-Rep Overview

The medial representation used is based on the medial axis of Blum [1]. In this framework, a geometric object is represented as a set of connected continuous medial manifolds. For 3D objects these medial manifolds are formed by the centers of all spheres that are interior to the object and tangent to the object's boundary at two or more points. The medial description is defined by the centers of the inscribed spheres and by the associated scalar field of their radii. Each continuous segment of the medial manifold represents a medial figure. In this paper we focus on 3D objects that can be represented by a single medial figure.

We sample the medial manifold $\mathcal{M}$ over a spatially regular lattice. Each sample point also includes first derivative information of the medial position and radius. The elements of this lattice are called *medial atoms*. A medial atom (Fig. 1) is defined as a 4-tuple $\mathbf{m} = \{\mathbf{x}, r, \mathbf{F}, \theta\}$, consisting of:

$\mathbf{x} \in \mathbb{R}^3$, the center of the inscribed sphere, $r \in \mathbb{R}^+$, the local width defined as the radius of the sphere, $\mathbf{F} \in \mathbf{SO}(3)$ an orthonormal local frame parameterized by $(\mathbf{b}, \mathbf{b}^\perp, \mathbf{n})$, where $\mathbf{n}$ is the normal to the medial manifold, $\mathbf{b}$ is the direction in the tangent plane of the fastest narrowing of the implied boundary sections, and $\theta \in [0, \pi)$ the object angle determining the angulation of the implied sections of boundary relative to $\mathbf{b}$. The medial atom implies two opposing boundary points, $\mathbf{y}_0, \mathbf{y}_1$, with respective boundary normals, $\mathbf{n}_0, \mathbf{n}_1$, which are given by

$$\mathbf{n}_0 = \cos(\theta)\mathbf{b} - \sin(\theta)\mathbf{n}, \qquad \mathbf{n}_1 = \cos(\theta)\mathbf{b} + \sin(\theta)\mathbf{n},$$
$$\mathbf{y}_0 = \mathbf{x} + r\mathbf{n}_0, \qquad \mathbf{y}_1 = \mathbf{x} + r\mathbf{n}_1.$$

For three dimensional slab-like figures (Fig. 1) the lattice of medial atoms is a quadrilateral mesh $\mathbf{m}_{ij}, (i, j) \in [1, m] \times [1, n]$. The sampling density of medial atoms in a lattice is inversely proportional to the radius of the medial description. Given an m-rep figure, we fit a smooth boundary surface to the model. We use a subdivision surface method [20] that interpolates the boundary positions and normals implied by each atom.

1.2. Lie Groups

Here we present a brief overview of Lie groups. For a detailed treatment see [7]. A Lie group G is a differentiable manifold that also forms an algebraic group, where the two group operations,

$$\mu : (x, y) \mapsto xy \quad : \quad G \times G \to G \quad \textit{Multiplication,}$$
$$\iota : x \mapsto x^{-1} \quad : \quad G \to G \quad \textit{Inverse,}$$

are differentiable mappings.

A Lie algebra $\mathfrak{g}$ is a vector space with a bilinear product $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$, called the Lie bracket, satisfying

$$[X, Y] = -[Y, X] \quad \textit{Antisymmetry,}$$
$$[[X, Y], Z] = [X, [Y, Z]] - [Y, [X, Z]] \quad \textit{Jacobi Identity,}$$

for all $X, Y, Z \in \mathfrak{g}$. Let e denote the identity element of a Lie group G. The tangent space at e, T_eG, forms a Lie algebra, which we will denote by $\mathfrak{g}$.

The exponential map, $\exp : \mathfrak{g} \to G$, provides a method for mapping vectors in the tangent space T_eG into G. Given a vector $\mathbf{v} \in \mathfrak{g}$, the point $\exp(\mathbf{v}) \in G$ is obtained by flowing to time 1 along the unique one-parameter subgroup emanating from e with initial velocity vector $\mathbf{v}$. When the Lie group is given a compatible Riemannian metric, this one-parameter subgroup is the unique geodesic at e with velocity $\mathbf{v}$. The exponential map is a diffeomorphism of a neighborhood of 0 in $\mathfrak{g}$ with a neighborhood of e in G. The inverse of the exponential map is called the log map. The geodesic distance between two points $g, h \in G$ is given by $||\log(g^{-1}h)||$.

For real numbers x and y the exponential map satisfies the identity $\exp(x)\exp(y) = \exp(x+y)$. This identity does not hold for noncommutative Lie groups, such as $\mathbf{SO}(3)$. In general define the *product in logarithmic coordinates* as a mapping $\mu : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$ such that

$$\exp(X)\exp(Y) = \exp(\mu(X, Y)).$$

A Taylor series expansion for μ about the point $(0, 0)$ is given by the Cambell-Baker-Hausdorff (CBH) formula:

$$\mu(X, Y) = X + Y + \frac{1}{2}[X, Y] + \mathcal{O}(|(X, Y)|^3). \qquad (1)$$

1.3. Discrete M-Rep as a Point on a Lie Group

In this section we show that the set of medial atoms is a Lie group. A medial atom's position is an element of $\mathbb{R}^3$, which is a standard Lie group under vector addition. The radius parameter is an element of the multiplicative Lie group of positive reals. The medial atom's frame is a 3D rotation, and the object angle is a 2D rotation. Both $\mathbf{SO}(2)$ and $\mathbf{SO}(3)$ are Lie groups under the composition of rotations. Thus, the set of all medial atoms forms a group $M = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbf{SO}(3) \times \mathbf{SO}(2)$, which we call the *medial group*. Since M is the direct product of four Lie groups, it also is a Lie group. Strictly speaking, the group M is a transformation group that *acts* on medial atoms. However, any medial atom may be represented in M as the transformation of a fixed identity atom, that is, the atom centered at the origin, with standard coordinate frame, radius 1, and object angle 0.

Now consider the set of m-rep models that consist of a $m \times n$ grid of medial atoms. These models form the space M^{mn}. Since this is simply the direct product of mn copies of M, it is a Lie group. Again, this is really a transformation group acting on the identity model, that is, the model consisting of an $m \times n$ lattice of identity atoms. Now, given the medial descriptions of a population of objects, we may consider each geometric model as a point on the Lie group M^{mn}.

In the definition of M above, one might think to replace the direct product $\mathbb{R}^3 \times \mathbb{R}^+ \times \mathbf{SO}(3)$ by the group of similarity transformations $\mathbf{Sim}(3) = \mathbb{R}^+ \times \mathbf{SO}(3) \ltimes \mathbb{R}^3$, where $\ltimes$ denotes a semidirect product. However, the semidirect product space does not have a Riemannian structure that is compatible with the algebraic structure, i.e., the Lie group exponential map does not give geodesics [22]. Since geodesic distance plays a key role in our work, we use the direct product space, which does have a compatible Riemannian metric.

1.4. Matrix Groups

The most common examples of Lie groups, and those which have the greatest application to computer vision, are the matrix groups [6]. These are all subgroups of the general linear group $\mathbf{GL}(n,\mathbb{R})$, the group of nonsingular $n \times n$ real matrices. The Lie algebra associated with $\mathbf{GL}(n,\mathbb{R})$ is $\mathbf{L}(\mathbb{R}^n,\mathbb{R}^n)$, the set of all $n \times n$ real matrices. The exponential map of a matrix $X \in \mathbf{L}(\mathbb{R}^n,\mathbb{R}^n)$ is the standard matrix exponent defined by the infinite series

$$\exp(X) = \sum_{k=0}^{\infty} \frac{1}{k!} X^k. \tag{2}$$

It is well-known that the rotation groups $\mathbf{SO}(2)$ and $\mathbf{SO}(3)$ are matrix subgroups of $\mathbf{GL}(2,\mathbb{R})$ and $\mathbf{GL}(3,\mathbb{R})$, respectively. The group of 3D rigid motions, $\mathbf{SE}(3)$, has also been well studied [17]. Related work includes the statistical analysis of directional data [15] and the study of shape spaces as complex projective spaces [14].

The 2D rotation group, $\mathbf{SO}(2)$, has corresponding Lie algebra $\mathfrak{so}(2)$, the set of 2×2 skew-symmetric matrices. Likewise, the Lie algebra for the 3D rotation group, $\mathbf{SO}(3)$, is the set of 3×3 skew-symmetric matrices, $\mathfrak{so}(3)$. We will use the notation

$$\mathbf{A}_\theta = \begin{pmatrix} 0 & -\theta \\ \theta & 0 \end{pmatrix}, \quad \mathbf{A}_\mathbf{v} = \begin{pmatrix} 0 & -v_1 & v_2 \\ v_1 & 0 & -v_3 \\ -v_2 & v_3 & 0 \end{pmatrix},$$

for elements of $\mathfrak{so}(2)$ and $\mathfrak{so}(3)$, respectively, where $\theta \in [0, 2\pi)$, and $\mathbf{v} = (v_1, v_2, v_3) \in \mathbb{R}^3$. Here, θ represents the angle of rotation in the plane. For 3D rotations the normalized vector $\bar{\mathbf{v}} = \frac{\mathbf{v}}{||\mathbf{v}||}$ is an axis of rotation, and the angle of rotation about that axis is $||\mathbf{v}||$.

The exponential map for $\mathfrak{so}(2)$ takes the form $\exp(\mathbf{A}_\theta) = \mathbf{R}_\theta$, where $\mathbf{R}_\theta$ is the matrix for a 2D rotation by θ. The exponential map for $\mathfrak{so}(3)$ is given by Rodrigues' formula [16]

$$\exp(\mathbf{A}_\mathbf{v}) = \begin{cases} \mathbf{I}_3, & \theta = 0, \\ \mathbf{I}_3 + \dfrac{\sin\theta}{\theta}\mathbf{A}_\mathbf{v} + \dfrac{1-\cos\theta}{\theta^2}{\mathbf{A}_\mathbf{v}}^2, & \theta \in (0,\pi), \end{cases} \tag{3}$$

where $\theta = \sqrt{\frac{1}{2}\mathrm{tr}({\mathbf{A}_\mathbf{v}}^{\mathrm{T}}\mathbf{A}_\mathbf{v})} = ||\mathbf{v}||$ in $[0,\pi)$.

Also, the logarithm for a matrix $\mathbf{R} \in \mathbf{SO}(3)$ is the matrix in $\mathfrak{so}(3)$ given by

$$\log(\mathbf{R}) = \begin{cases} \mathbf{0}, & \theta = 0, \\ \dfrac{\theta}{2\sin\theta}(\mathbf{R} - \mathbf{R}^T), & |\theta| \in (0,\pi), \end{cases} \tag{4}$$

where θ satisfies $\mathrm{tr}(\mathbf{R}) = 2\cos\theta + 1$.

1.5. The Exponential and Log Maps for M-reps

Now we are ready to define the exponential and log maps for the medial group M. The Lie algebra of M is the product space $\mathfrak{m} = \mathbb{R}^3 \times \mathbb{R} \times \mathfrak{so}(3) \times \mathfrak{so}(2)$. We define the norm of a vector $\mathbf{u} = (\mathbf{x}, \rho, \mathbf{A}_\mathbf{v}, \mathbf{A}_\theta) \in \mathfrak{m}$ as

$$||\mathbf{u}|| = \left(||\mathbf{x}||^2 + \rho^2 + \frac{1}{2}||\mathbf{A}_\mathbf{v}||^2 + \frac{1}{2}||\mathbf{A}_\theta||^2 \right)^{\frac{1}{2}},$$

where the matrix norms are Frobenius norms.

The exponential map for $\mathbb{R}^3$ is the identity map, and the exponential map for $\mathbb{R}$ is the familiar real exponential function. Combined with the exponential maps for the rotation groups given above, the exponential map for the medial group M is

$$\begin{aligned} \exp &: \mathfrak{m} \to M \\ &: (\mathbf{x}, \rho, \mathbf{A}_\mathbf{v}, \mathbf{A}_\theta) \mapsto (\mathbf{x}, \mathrm{e}^\rho, \exp(\mathbf{A}_\mathbf{v}), \exp(\mathbf{A}_\theta)), \end{aligned}$$

where we have abused notation by reusing $\exp$, but it is clear which exponential map we mean by the context. The corresponding log map is

$$\begin{aligned} \log &: M \to \mathfrak{m} \\ &: (\mathbf{x}, r, \mathbf{F}, \mathbf{R}_\theta) \mapsto (\mathbf{x}, \log(r), \log(\mathbf{F}), \log(\mathbf{R}_\theta)). \end{aligned}$$

2. Means in Lie Groups

In this section we formulate two different notions of means on Lie groups. We then present a method for computing the mean of a collection of m-rep models in the Lie group M^n.

2.1. Intrinsic vs. Extrinsic Means

Given a set of points $x_1, \ldots, x_n \in \mathbb{R}^d$, the arithmetic mean $\bar{x} = \frac{1}{n}\sum_{i=1}^n x_i$ is the point that minimizes the sum-of-squared Euclidean distances to the given points, i.e.,

$$\bar{x} = \arg\min_{x \in \mathbb{R}^d} \sum_{i=1}^n ||x - x_i||^2.$$

Since a general Lie group G may not form a vector space, the notion of an additive mean is not necessarily valid. However, like the Euclidean case, the mean of a set of points on G can be formulated as the point which minimizes the sum-of-squared distances to the given points. This formulation depends on the definition of distance. One way to define distance on G is to embed it in a Euclidean space and use the induced Euclidean distance. This notion of distance is extrinsic to G, that is, it depends on the ambient space and the choice of embedding. Given an embedding

$\Phi : G \to \mathbb{R}^d$, define the *extrinsic mean* [9] of a collection of points $x_1, \ldots, x_n \in G$ as

$$\mu_\Phi = \arg\min_{x \in G} \sum_{i=1}^{n} ||\Phi(x) - \Phi(x_i)||^2.$$

Given the above embedding of G, we can also compute the arithmetic (Euclidean) mean of the embedded points and then project this mean onto the manifold G. This projected mean is equivalent to the above definition of the extrinsic mean (see [18]). Define a projection mapping $\pi : \mathbb{R}^d \to G$ as

$$\pi(x) = \arg\min_{s \in G} ||\Phi(s) - x||^2.$$

Then the extrinsic mean is also given by

$$\mu_\Phi = \pi\Big(\frac{1}{n}\sum_{i=1}^{n} \Phi(x_i)\Big).$$

A more natural choice of distance is the Riemannian distance on G. The Riemannian distance between two points is the length of the shortest geodesic curve between the points. This definition of distance depends only on the intrinsic geometry of G. We now define the *intrinsic mean* of a collection of points $x_1, \ldots, x_n \in G$ as the minimizer in G of the sum-of-squared Riemannian distances to each point. Thus the intrinsic mean is

$$\mu = \arg\min_{x \in G} \sum_{i=1}^{n} d(x, x_i)^2,$$

where $d(\cdot, \cdot)$ denotes Riemannian distance on G.

2.2. M-rep Averages

Intrinsic means in $\mathbf{SO}(3)$ have been well studied [2, 16]. Here we extend these ideas to develop an algorithm for computing the intrinsic mean of a collection of m-rep models, that is, an intrinsic mean of points on the Lie group M^n. The algorithm we present proceeds by iteratively computing first-order approximations to the true mean. This algorithm is similar in nature to the technique found in [2] for computing averages of spherical data. Our method is a generalization to averages in M^n.

The Riemannian distance between two points $\mathbf{M}_1, \mathbf{M}_2 \in M^n$ is given by

$$d(\mathbf{M}_1, \mathbf{M}_2) = ||\log(\mathbf{M}_1^{-1}\mathbf{M}_2)||. \tag{5}$$

Thus the intrinsic mean of a collection of points $\mathbf{M}_1, \ldots, \mathbf{M}_N$ is

$$\mu = \arg\min_{\mathbf{M} \in M^n} \sum_{i=1}^{n} ||\log(\mathbf{M}_i^{-1}\mathbf{M})||^2.$$

Using the CBH formula (1) up to first-order terms, we may approximate the Riemannian distance in (5) with the product in logarithmic coordinates as

$$d(\mathbf{M}_1, \mathbf{M}_2) \approx ||\log(\mathbf{M}_2) - \log(\mathbf{M}_1)||. \tag{6}$$

Now the sum-of-squared approximated distances is minimized by the arithmetic mean of the log of the points. The first-order approximation to the intrinsic mean is given by

$$\hat{\mu} = \exp\Big(\frac{1}{n}\sum_{i=1}^{n} \log(\mathbf{M}_i)\Big)$$

According to the CBH formula, $\hat{\mu}$ is a first-order approximation of the intrinsic mean. Also, the error in this approximation is larger when the points are far from the identity. Thus we left-multiply all points by $\hat{\mu}^{-1}$ so that $\hat{\mu}$ is moved to the identity. Now we compute the mean of these residual points and combine this with $\hat{\mu}$ to arrive at a new approximation to the mean. This process is repeated until the mean of the residuals is sufficiently near the identity. Summarizing, we have

Algorithm 1:
Input: $\mathbf{M}_1, \ldots, \mathbf{M}_n \in M^n$, m-rep models
Output: $\mu \in M^n$, the intrinsic mean
  $\mu = \mathbf{M}_1$
  Do
    $\Delta\mathbf{M}_i = \mu^{-1}\mathbf{M}_i$
    $\Delta\mu = \exp(\frac{1}{n}\sum_{i=1}^{n} \log(\Delta\mathbf{M}_i))$
    $\mu = \mu\Delta\mu$
  While $||\log(\Delta\mu)|| > \epsilon$.

Following the argument in [2], this is in fact a gradient descent method, and the true intrinsic mean is a stable point of the algorithm. Notice that each of the $\mathbb{R}^3, \mathbb{R}^+$, and $\mathbf{SO}(2)$ components in M^n converge in a single iteration, since they are commutative groups. Figure 2 shows twelve m-rep models of human kidneys and the resulting intrinsic mean of the models, computed by the above algorithm.

3. Principal Geodesic Curves

We begin by reviewing PCA in a linear space. In a vector space V, each unit vector $\mathbf{v}$ determines a one-dimensional subspace $S_\mathbf{v} = \{t\mathbf{v} : t \in \mathbb{R}\}$. Let $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_m \in V$. Without loss of generality, assume the mean vector of $\{\mathbf{v}_i\}$ is $\mathbf{0}$. The first principal component can then be characterized as the one-dimensional subspace $S_{\mathbf{u}^{(1)}}$ where $\mathbf{u}^{(1)}$ satisfies

$$\mathbf{u}^{(1)} = \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} \|\mathbf{v}_i - (\mathbf{v}_i \cdot \mathbf{v})\mathbf{v}\|^2.$$

In other words, if $d_i(\mathbf{v})$ is the distance from $\mathbf{v}_i$ to $S_\mathbf{v}$, then the first principal direction is defined by the vector $\mathbf{v}$ such

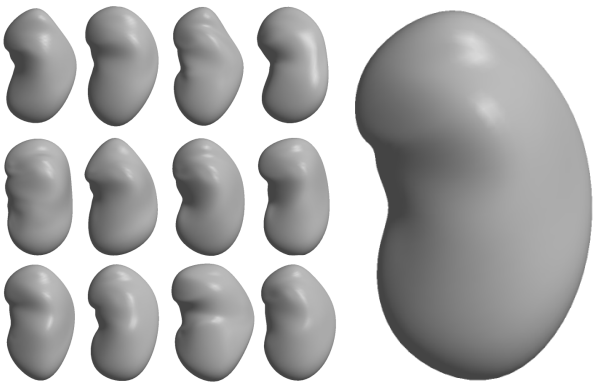

Figure 2. Boundary surface display of the 12 original m-rep kidney models (left). The intrinsic mean of the kidney models (right).

that $\sum_{i=1}^{m} d_i^2(\mathbf{v})$ is minimized. The eigenvalue associated with the first component is simply given by

$$\lambda_1 = \frac{1}{m}\sum_{i=1}^{m}(\mathbf{v}_i \cdot \mathbf{u}^{(1)})^2,$$

where $\mathbf{v}_i \cdot \mathbf{u}^{(1)}$ is the projection of $\mathbf{v}_i$ onto the linear subspace $S_{\mathbf{u}^{(1)}}$. Similarly, for $k > 1$, the k-th principal direction is defined recursively by the vector

$$\mathbf{u}^{(k)} = \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} \Big\| \mathbf{v}_i - \sum_{l=1}^{k-1} (\mathbf{v}_i \cdot \mathbf{u}^{(l)})\mathbf{u}^{(l)} - (\mathbf{v}_i \cdot \mathbf{v})\mathbf{v} \Big\|^2.$$

In this section, we generalize these concepts to finite dimensional Lie groups.

Let G be a Lie group. The corresponding Lie algebra $\mathfrak{g}$ is a vector space. For an arbitrary unit vector $\mathbf{v} \in \mathfrak{g}$, we can define a one-parameter subgroup $H_{\mathbf{v}}$ of G as

$$H_{\mathbf{v}} \triangleq \{\exp(t\mathbf{v}) \in G : t \in \mathbb{R}\}, \tag{7}$$

where $\exp(\cdot)$ is the exponential map. It is easy to see that $H_{\mathbf{v}}$ is a geodesic curve in G whose tangent direction at the identity is given by $\mathbf{v}$. For any $g \in G$, the distance from g to $H_{\mathbf{v}}$ is defined as

$$d(g, H_{\mathbf{v}}) \triangleq \min_t d\big(g, \exp(t\mathbf{v})\big), \tag{8}$$

where the distance on the right hand side is the geodesic distance on G. Among all values of t such that (8) is minimized, let t^* to be the one that has the smallest absolute value. The *projection* of g on $H_{\mathbf{v}}$ is defined to be $\exp(t^*\mathbf{v})$.

Suppose that $g_1, g_2, \ldots, g_m$ are m elements of a Lie group G, and that the intrinsic mean of them is μ. The *first principal geodesic curve* for these elements is defined as the one parameter subgroup $H_{\mathbf{u}^{(1)}}$ of G, where

$$\mathbf{u}^{(1)} = \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} d^2(\mu^{-1}g_i, H_{\mathbf{v}}). \tag{9}$$

Let $p_{i,1}$ be the projection of $\mu^{-1}g_i$ on $H_{\mathbf{u}^{(1)}}$, and define $g_i^{(1)} = p_{i,1}^{-1}\mu^{-1}g_i$. For $k \geq 2$, the k-th principal geodesic curve is defined by the vector $\mathbf{u}^{(k)}$ which is given recursively by

$$\begin{aligned} \mathbf{u}^{(k)} &= \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} d^2\Big(g_i^{(k-1)}, H_{\mathbf{v}}\Big), \\ g_i^{(k)} &= p_{i,k}^{-1} g_i^{(k-1)}, \end{aligned} \tag{10}$$

where $p_{i,k}$ is the projection of $g_i^{(k-1)}$ on $H_{\mathbf{u}^{(k)}}$.

Notice that the first k principal curves yield a decomposition for $\{g_i\}$:

$$g_i = \mu p_{i,1} p_{i,2} \cdots p_{i,k} g_i^{(k)}, \quad k \geq 1, i = 1, \ldots, m.$$

Analogous to the eigenvalues in the linear case, we define the energy associated with the k-th principal geodesic to be

$$\lambda_k = \frac{1}{m}\sum_{i=1}^{m} t_{i,k}^2,$$

where $t_{i,k}$ is the geodesic distance between $p_{i,k}$ and the identity.

4. Approximating Principal Geodesic Curves

Given the ability to compute geodesic distances on G, principal geodesic analysis (PGA) reduces to a nonlinear optimization problem defined by (9) and (10). This optimization can be impractical in high dimensions. For efficient computation, we now show that the principal geodesic curves can be approximated via a linear PCA in the tangent space to the mean, $T_\mu G$. Using the CBH formula and the approximation (6), equation (9) can be approximated as follows

$$\begin{aligned} \mathbf{u}^{(1)} &= \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} \min_t \; d^2(\mu^{-1}g_i, \exp(t\mathbf{v})) \\ &\approx \arg\min_{\|\mathbf{v}\|=1} \sum_{i=1}^{m} \min_t \| \log(\mu^{-1}g_i) - t\mathbf{v}\|^2. \end{aligned}$$

Noticing that $\mathbf{v}_i = \log(\mu^{-1}g_i), i = 1, \ldots, m$, are vectors in the tangent space $T_\mu G$, the above minimization problem becomes the standard principal component analysis in $T_\mu G$.

The geodesic curves defined via the exponential mapping $H_{\mathbf{u}^{(k)}} = \exp(t\mathbf{u}^{(k)})$ are the approximations of the principal geodesic curves defined in the previous section.

More precisely, for m-rep models we compute the principal geodesic analysis by

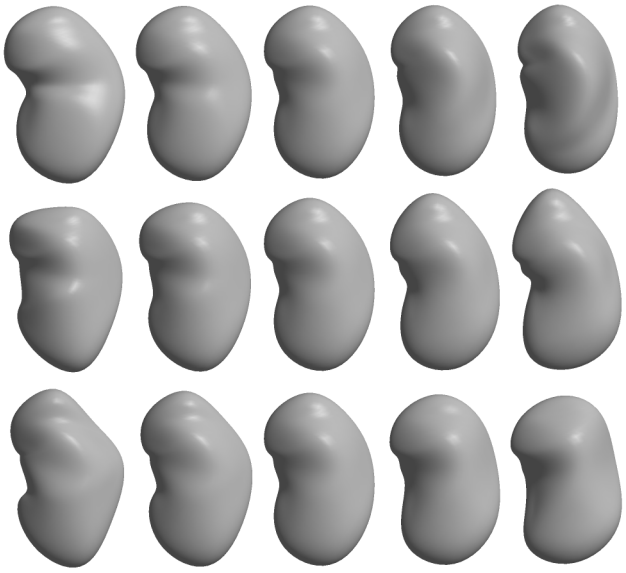

Figure 3. The first three modes of deformation for kidney shapes. Each row displays the models corresponding to $\{-3\sqrt{\lambda_i}, -1.5\sqrt{\lambda_i}, 0, +1.5\sqrt{\lambda_i}, +3\sqrt{\lambda_i}\}$ along the i^{th} principal geodesic.

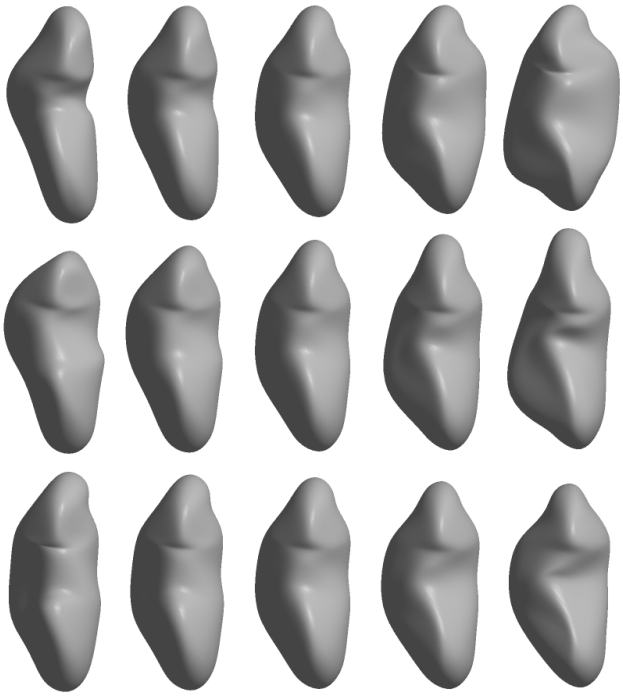

Figure 4. The same modes of deformation as Fig. 3 shown from the side.

Algorithm 2:
Input: M-rep models, $\mathbf{M}_1, \ldots, \mathbf{M}_n \in M^n$
Output: Principal directions, $\mathbf{u}^{(k)} \in T_\mu M^n$
Variances, $\lambda_k \in \mathbb{R}$
$\mu =$ intrinsic mean of $\{\mathbf{M}_i\}$ (Algorithm 1)
$\mathbf{x}_i = \log(\mu^{-1}\mathbf{M}_i)$
$\mathbf{S} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^T$
$\{\mathbf{u}^{(k)}, \lambda_k\}$ = eigenvectors/eigenvalues of $\mathbf{S}$.

Analogous to linear PCA models, we may choose a subset of the principal directions $\mathbf{u}^{(k)}$ that is sufficient to describe the variability of the m-rep shape space. New m-rep models may be generated within this subspace of typical objects. Given a set of coefficients $\{\alpha_1, \ldots, \alpha_l\}$, we generate a new m-rep model by

$$\mathbf{M} = \mu \exp\Big(\sum_{k=1}^{l} \alpha_k \mathbf{u}^{(k)}\Big),$$

where α_k is chosen to be within $[-3\sqrt{\lambda_k}, 3\sqrt{\lambda_k}]$.

5. Results

In this section we present the results of applying our Lie group PGA method to a population of medial models of the human kidney (Fig. 2). The kidney m-rep models were automatically generated by the method described in [19], which chooses the medial topology and sampling that is sufficient to represent the population of objects. Twelve models were fit to hand segmentations of the kidneys in CT data. The sampling on each m-rep model was 3×5, and thus each model was a point on the Lie group M^{15}.

A coarsely sampled boundary surface was obtained for each model by taking the boundary points implied by each medial atom. The kidneys were aligned using a generalized Procrustes analysis [8] on these boundary points. The mean object (Fig. 2) was generated as the mean of the twelve aligned m-rep models, using the intrinsic mean computation on M^{15} described in §2.2. A linear PCA was performed in the tangent space to the intrinsic mean, as described in §4, producing a PGA shape space. The first three modes are shown in Figures 3 and 4.

6. Discussion

We present a new approach to describing shape variability through principal geodesic analysis of medial representations. While m-rep parameters are not linear vector spaces, we show that they are indeed Lie groups. We develop methods for computing averages and principal geodesic analyses of m-reps.

We point out that there is a method called principal

curves [12], which has a similar name to principal geodesic curves. However, the two methods are only loosely related. Principal curves are smooth curves that are fit to data in Euclidean space by minimizing the sum-of-squared Euclidean distances to the data. Principal geodesic analysis on the other hand concerns data that lie in a Lie group, rather than Euclidean space. Principal geodesic curves are intrinsic to the underlying space, and they minimize sum-of-squared geodesic distances in that space.

We expect that the methods presented in this paper will have application beyond m-reps. Lie group PGA is a promising technique for describing the variability of data that is inherently nonlinear. Statistics on linear models may benefit from the addition of nonlinear information. For instance the point distribution model [4] might be augmented with surface normals, represented as orientations, and handled under the Lie group framework.

We plan to extend our analysis to more complex m-rep models. This includes objects consisting of several figures, i.e., objects that have a branched medial axis. Also, we intend to handle scenes containing multiple objects.

Acknowledgments

We would like to thank Dr. Steve Pizer for his advice and for providing the driving problem of this work. We would also like to thank Manjari Rao, Gregg Tracton, and Dr. Ed Chaney for providing the data used in the results. We thank Dr. Ulf Grenander and Dr. Anuj Srivastava for useful discussions on Lie groups. This work was done with support from NCI grant P01 CA47982.

References

[1] H. Blum and R. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167–180, 1978.

[2] S. R. Buss and J. P. Fillmore. Spherical averages and applications to spherical splines and interpolation. *ACM Transactions on Graphics*, 20(2):95–126, 2001.

[3] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *Fifth European Conference on Computer Vision*, pages 484–498, 1998.

[4] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[5] J. Csernansky, S. Joshi, L. Wang, J. Haller, M. Gado, J. Miller, U. Grenander, and M. Miller. Hippocampal morphometry in schizophrenia via high dimensional brain mapping. In *Proceedings National Academy of Sciences*, pages 11406–11411, 1998.

[6] M. L. Curtis. *Matrix Groups*. Springer-Verlag, 1984.

[7] J. J. Duistermaat and J. A. C. Kolk. *Lie Groups*. Springer, 2000.

[8] C. Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society*, 53(2):285–339, 1991.

[9] U. Grenander, M. I. Miller, and A. Srivastava. Hilbert-Schmidt lower bounds for estimators on matrix Lie groups for ATR. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):790–802, 1998.

[10] I. T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 1986.

[11] S. Joshi, S. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, and J. S. Marron. Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *Transactions on Medical Imaging*, 21(5), 2002.

[12] B. Kégl, A. Krzyzak, T. Linder, and K. Zeger. Learning and design of principal curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3), 2000.

[13] A. Kelemen, G. Szekely, and G. Gerig. Three-dimensional model-based segmentation. *Transactions on Medical Imaging*, 18(10):828–839, 1999.

[14] D. G. Kendall. Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16:18–121, 1984.

[15] K. V. Mardia. *Directional Statistics*. John Wiley and Sons, 1999.

[16] M. Moakher. Means and averaging in the group of rotations. *SIAM Journal on Matrix Analysis and Applications*, 24(1):1–16, 2002.

[17] R. Murray, Z. Li, and S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[18] A. Srivastava and E. Klassen. Monte-Carlo extrinsic estimators of manifold-valued parameters. *IEEE Transactions on Signal Processing*, 50(2):299–308, 2001.

[19] M. Styner and G. Gerig. Medial models incorporating object variability for 3D shape analysis. In *Information Processing in Medical Imaging*, pages 502–516, 2001.

[20] A. Thall. Fast C^2 interpolating subdivision surfaces using iterative inversion of stationary subdivision rules. Technical report, University of North Carolina Department of Computer Science, 2002. http://midag.cs.unc.edu/pub/papers/Thall_TR02-001.pdf.

[21] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal for Cognitive Neuroscience*, 3(1):71–86, 1991.

[22] M. Zefran, V. Kumar, and C. Croke. Metrics and connections for rigid body kinematics. *The International Journal of Robotics Research*, 18(2), 1999.