Gaussian Distributions on Lie Groups and Their Application to Statistical Shape Analysis

P. Thomas Fletcher, Sarang Joshi, Conglin Lu, Stephen Pizer

Medical Image Display and Analysis Group,
University of North Carolina at Chapel Hill
fletcher@cs.unc.edu

Abstract. The Gaussian distribution is the basis for many methods used in the statistical analysis of shape. One such method is principal component analysis, which has proven to be a powerful technique for describing the geometric variability of a population of objects. The Gaussian framework is well understood when the data being studied are elements of a Euclidean vector space. This is the case for geometric objects that are described by landmarks or dense collections of boundary points. We have been using medial representations, or m-reps, for modelling the geometry of anatomical objects. The medial parameters are not elements of a Euclidean space, and thus standard PCA is not applicable. In our previous work we have shown that the m-rep model parameters are instead elements of a Lie group. In this paper we develop the notion of a Gaussian distribution on this Lie group. We then derive the maximum likelihood estimates of the mean and the covariance of this distribution. Analogous to principal component analysis of covariance in Euclidean spaces, we define principal geodesic analysis on Lie groups for the study of anatomical variability in medially-defined objects. Results of applying this framework on a population of hippocampi in a schizophrenia study are presented.

1 Introduction

Shape analysis is emerging as an important area of image processing and computer vision. Model-based approaches [1, 2, 3] are popular due to their ability to robustly represent objects found in images. Principal component analysis (PCA) [4] 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 [6] 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 Gaussian distributions on this Lie group and derive the maximum likelihood estimates (MLEs) of the mean and covariance. Using these distributions, we introduce principal geodesic analysis (PGA), the extension of PCA to Lie groups. We apply this framework to the statistical analysis of shape using medial representations. As the medial representation is fundamental to our analysis, we describe it briefly.

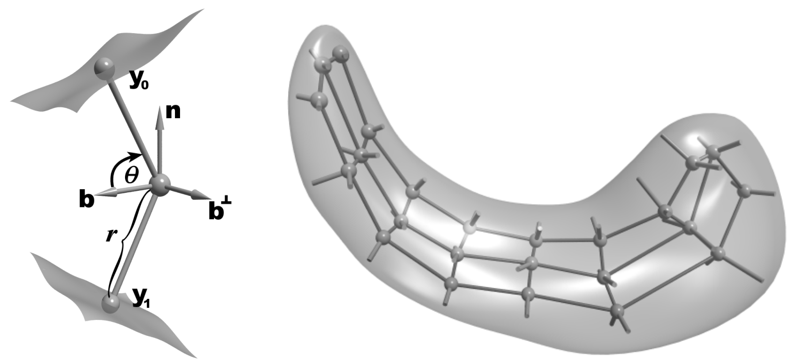

Fig. 1. Medial atom with a cross-section of the boundary surface it implies (left). An m-rep model of a hippocampus and its boundary surface (right).

1.1 M-Rep Overview

The medial representation used is based on the medial axis of Blum [7]. In this framework, a 3D geometric object is represented as a set of connected continuous medial manifolds, which 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. 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

$$\begin{aligned} \mathbf{n}_0 &= \cos(\theta)\mathbf{b} - \sin(\theta)\mathbf{n}, & \mathbf{n}_1 &= \cos(\theta)\mathbf{b} + \sin(\theta)\mathbf{n}, \\ \mathbf{y}_0 &= \mathbf{x} + r\mathbf{n}_0, & \mathbf{y}_1 &= \mathbf{x} + r\mathbf{n}_1. \end{aligned} \tag{1}$$

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 [8] 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 [9]. A Lie group G is a differentiable manifold that also forms an algebraic group, where the two group operations,

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

are differentiable mappings.

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 geodesic emanating from e with initial velocity vector $\mathbf{v}$. 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 $|| \log(g^{-1}h)||$.

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

We now show that a set of medial atoms defining an m-rep object can be represented as a point on 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.

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. 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}.

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 [10]. 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 [11]. Related work includes the statistical analysis of directional data [12] and the study of shape spaces as complex projective spaces [13].

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 [14]

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

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

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

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 $\mathfrak{m} = \mathbb{R}^3 \times \mathbb{R} \times \mathfrak{so}(3) \times \mathfrak{so}(2)$. 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 Gaussian Distributions on Lie Groups

In this section we develop Gaussian distributions on M^n, the Lie group of m-rep figures with n atoms. We begin by developing Gaussian distributions on each of the factors in the product space $M = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbf{SO}(3) \times \mathbf{SO}(2)$. We define Gaussian distributions on Lie groups following Grenander [15]. A Gaussian distribution on a Lie group with mean at the identity element is a solution to the heat equation defined in the local coordinates of the Lie group:

$$\begin{aligned}\frac{\partial f}{\partial t} &= \Delta f = \mathrm{div}(\mathrm{grad} f)\\ &= g^{ij}\Big(\frac{\partial^2 f}{\partial x^i \partial x^j} - \Gamma^k_{ij}\frac{\partial f}{\partial x^k}\Big),\end{aligned}$$

where g^{ij} are the components of the inverse of the Riemannian metric, and Γ^k_{ij} are the Christoffel symbols [16]. Indeed, the Gaussian distribution in $\mathbb{R}^n$, given by the density

$$p(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n|\mathbf{\Sigma}|}} \exp\Big(- \frac{(\mathbf{x}-\boldsymbol{\mu})^T \mathbf{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}{2}\Big), \tag{5}$$

is the solution of the heat equation in $\mathbb{R}^n$. The case $n = 3$ gives the Gaussian distribution for medial atom positions.

2.1 Gaussian Distributions on $\mathbb{R}^+$

For the Lie group of positive reals under multiplication, local coordinates are given by the logarithm. The solution to the heat equation on $\mathbb{R}^+$ is given by the lognormal density:

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp\Big(- \frac{(\log x - \log \mu)^2}{2\sigma^2}\Big). \tag{6}$$

Given samples $x_1, \ldots, x_N \in \mathbb{R}^+$ that are independently distributed by the lognormal distribution, the maximum likelihood estimates for the mean and variance are

$$\hat{\mu} = \Big(\prod_{i=1}^N x_i\Big)^{\frac{1}{N}}, \qquad \hat{\sigma}^2 = \frac{1}{N}\sum_{i=1}^N (\log x_i - \log \hat{\mu})^2.$$

Notice that $\hat{\mu}$, given by the geometric average, is the point that minimizes the sum-of-squared geodesic distances in $\mathbb{R}^+$, i.e., it minimizes $\sum_{i=1}^n \log(\mu^{-1}x_i)^2$.

2.2 Gaussian Distributions on $\mathbf{SO}(2)$

Consider the parametrization of $\mathbf{SO}(2)$ by the rotation angle $\theta \in [0, 2\pi)$. Notice that $\mathbf{SO}(2)$ is isomorphic to the unit circle S^1 via the mapping $\theta \mapsto \mathrm{e}^{\mathrm{i}\theta}$. The Gaussian distribution on $\mathbf{SO}(2)$ with mean μ and standard deviation σ is given by

$$p(\theta) = \frac{1}{\sqrt{2\pi}\sigma} \sum_{k=-\infty}^{\infty} \exp\Big(- \frac{(\theta - \mu - 2\pi k)^2}{2\sigma^2}\Big), \tag{7}$$

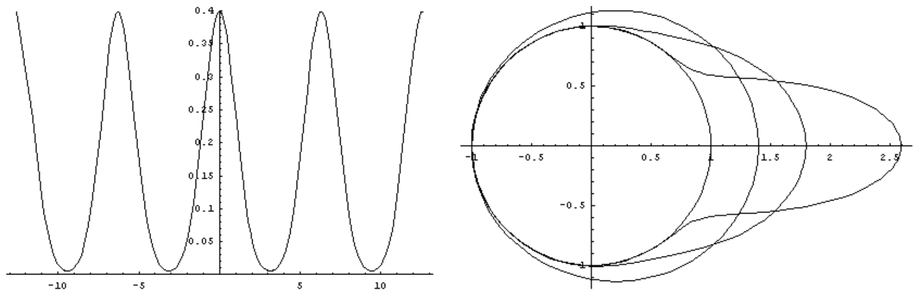

Fig. 2. Solution to the heat equation, cyclic on $[-\pi, \pi]$ (left). Wrapped Gaussian over the unit circle with $\sigma = 1$, 0.5, and 0.25 (right).

which solves the heat equation with cyclic boundary conditions on $[-\pi, \pi]$. This is sometimes known as the "wrapped Gaussian" on the circle (Fig. 2.2).

We now derive the maximum likelihood estimate for the mean and covariance, given samples $\theta_i \in [0, 2\pi), i = 1, \ldots, N$ that are independently distributed according to (7). To begin we assume that $\sigma = 1$, and we find the maximum likelihood estimate of the mean, which is given by

$$\hat{\mu} = \underset{\mu \in [0,2\pi)}{\arg\max} \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \exp\Big(- \frac{(\theta_i - \mu - 2\pi k)^2}{2} \Big). \tag{8}$$

Notice that since the quadratic exponential is an even function, its derivative is odd, and

$$\frac{\partial}{\partial \mu} \exp\Big(- \frac{(\theta_i - \mu - 2\pi k)^2}{2} \Big) = - \frac{\partial}{\partial \mu} \exp\Big(- \frac{(\theta_i - \mu + 2\pi k)^2}{2} \Big)$$

for a fixed integer k. Thus the derivative of the summation in (8) reduces to just the $k = 0$ term, and the maximization problem becomes

$$\hat{\mu} = \underset{\mu \in [0,2\pi)}{\arg\max} \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi}} \exp\Big(- \frac{(\theta_i - \mu)^2}{2} \Big). \tag{9}$$

This is just the equation for the maximum likelihood estimate of the mean for the Gaussian distribution. Therefore, we have $\hat{\mu} = \frac{1}{N} \sum_{i=1}^{N} \theta_i$. This equation can lead to ambiguities (see [12]), due to the multiple possible representations for the angles θ_i, e.g. we may take $\theta_i \in [0, 2\pi)$ or $\theta_i \in [-\pi, \pi)$. However, medial atom object angles always lie within $[0, \pi)$, and thus this ambiguity does not arise (see [17]).

For deriving the maximum likelihood estimate of variance, consider the log-likelihood

$$l(\sigma; \hat{\mu}, \theta_1, \ldots, \theta_N) = \sum_{i=1}^{N} \log\Big(\frac{1}{\sqrt{2\pi}\sigma} \Big) + \log\Bigg[\sum_{k=-\infty}^{\infty} \exp\Big(- \frac{(\theta_i - \hat{\mu} - 2\pi k)^2}{2\sigma^2} \Big) \Bigg].$$

Differentiation with respect to σ gives

$$\frac{\partial l}{\partial \sigma} = -\frac{N}{\sigma} + \frac{1}{\sigma^3} \sum_{i=1}^{N} (\theta_i - \hat{\mu})^2 + \frac{1}{\sigma^3} \sum_{i=1}^{N} \left[\frac{\sum_{k=-\infty}^{\infty} (2\pi k)^2 \exp\big(- \frac{(\theta_i - \hat{\mu} - 2\pi k)^2}{2\sigma^2} \big)}{\sum_{k=-\infty}^{\infty} \exp\big(- \frac{(\theta_i - \hat{\mu} - 2\pi k)^2}{2\sigma^2} \big)} \right].$$

The third term above, although it converges, does not yield a closed-form solution. However, according to Mardia [12], the wrapped Gaussian density (7) is well approximated by just the $k = 0$ term of the summation when $\sigma^2 \leq 2\pi$. This is certainly the case for medial atom object angles, which are tightly distributed. Thus we keep only the $k = 0$ term in the above summation, implying that

$$\frac{\partial l}{\partial \sigma} \approx -\frac{N}{\sigma} + \frac{1}{\sigma^3} \sum_{i=1}^{N} (\theta_i - \hat{\mu})^2.$$

Setting the above equation to zero and solving for σ, we get the approximated MLE of the variance as

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{i=1}^{N} (\theta_i - \hat{\mu})^2. \tag{10}$$

2.3 The Wrapped Gaussian Distribution on $\mathbf{SO}(3)$

Analogous to $\mathbf{SO}(2)$ we use the log map to define a wrapped Gaussian distribution on $\mathbf{SO}(3)$ with mean μ. We note that this density is not a solution of the heat equation on $\mathbf{SO}(3)$. Let $\mathbf{u}(x) = \Phi(\log(\mu^{-1}x)) \in \mathbb{R}^3$, where $\Phi : \mathfrak{so}(3) \to \mathbb{R}^3$, $\Phi(\mathbf{A_v}) = \mathbf{v}$, is the canonical isomorphism. Let $\bar{\mathbf{u}}(x) = \frac{\mathbf{u}(x)}{|\mathbf{u}(x)|}$. Following (7) the wrapped Gaussian density on $\mathbf{SO}(3)$ becomes

$$p(x) = \frac{1}{\sqrt{(2\pi)^3 |\mathbf{\Sigma}|}} \sum_{k=-\infty}^{\infty} \exp\Big(-\frac{1}{2} (\mathbf{u}(x) - 2\pi k \bar{\mathbf{u}}(x))^T \mathbf{\Sigma}^{-1} (\mathbf{u}(x) - 2\pi k \bar{\mathbf{u}}(x)) \Big). \tag{11}$$

Here $\mu \in \mathbf{SO}(3)$ is the mean rotation, and the covariance structure is defined as a quadratic form on the Lie algebra $\mathfrak{so}(3)$, represented as the 3×3 covariance matrix $\mathbf{\Sigma}$.

Given samples $x_1, \ldots, x_N \in \mathbf{SO}(3)$ independently distributed according to the density (11), we derive the maximum likelihood estimate for the mean and covariance. Focusing on the MLE of the mean, we may assume, without loss of generality, that the covariance is identity. The joint density is given by the product density

$$p(\mu; x_1, \ldots, x_N) = \prod_{i=1}^{N} \frac{1}{\sqrt{(2\pi)^3}} \sum_{k=-\infty}^{\infty} \exp\Big(-\frac{1}{2} ||\mathbf{u}(x_i) - 2\pi k \bar{\mathbf{u}}(x_i)||^2 \Big). \tag{12}$$

Notice that geodesics of $\mathbf{SO}(3)$ are isomorphic to $\mathbf{SO}(2)$, and the density (11) restricted to a geodesic reduces to the wrapped Gaussian on $\mathbf{SO}(2)$. Now we can use the same argument from the previous section to show that the derivatives in the summation cancel out. Therefore, maximizing $p(\mu; x_1, \ldots, x_N)$ is equivalent to maximizing

$$\prod_{i=1}^{N} \frac{1}{\sqrt{(2\pi)^3}} \exp\Big(-\frac{1}{2} ||\mathbf{u}(x_i)||^2 \Big). \tag{13}$$

Now, taking the log of the above, the MLE of the mean becomes

$$\hat{\mu} = \underset{\mu \in \mathbf{SO}(3)}{\arg\min} \sum_{i=1}^{N} ||\Phi(\log(\mu^{-1} x_i))||^2. \tag{14}$$

Notice that $||\Phi(\log(\mu^{-1}x_i))||$ is the Riemannian distance from μ to x_i. Hence, the MLE of the mean is also the point that minimizes the sum-of-squared geodesic distances to the samples. This is also referred to as the intrinsic mean on $\mathbf{SO}(3)$ [14]. An iterative algorithm for computing the intrinsic mean is given in [17].

As in the case for $\mathbf{SO}(2)$, we assume that the variance is sufficiently small, that is, $\lambda^2 \leq 2\pi$, where λ is an eigenvalue of $\mathbf{\Sigma}$. Using the same argument as in the previous section, the approximated maximum likelihood estimate of the covariance is

$$\hat{\mathbf{\Sigma}} = \frac{1}{N}\sum_{i=1}^{N}\Phi(\log(\hat{\mu}^{-1}x_i))\Phi(\log(\hat{\mu}^{-1}x_i))^T. \tag{15}$$

2.4 Gaussian Distributions on the Medial Group

We now combine the distributions developed on the factors $\mathbb{R}^3$, $\mathbb{R}^+$, $\mathbf{SO}(3)$, and $\mathbf{SO}(2)$ to define a Gaussian distribution on the Lie group M. As M is a direct product of these Lie groups, the Gaussian distribution on M is the product distribution given by

$$p(x) = \frac{1}{\sqrt{(2\pi)^8|\mathbf{\Sigma}|}}\sum_{k=-\infty}^{\infty}\exp\Big(-\frac{1}{2}(\mathbf{u}(x)-2\pi k\rho(\bar{\mathbf{u}}(x)))^T\mathbf{\Sigma}^{-1}(\mathbf{u}(x)-2\pi k\rho(\bar{\mathbf{u}}(x)))\Big).$$

Here $\mathbf{u}(x) = \log(\mu^{-1}x) \in \mathfrak{m}$ is represented as an 8-vector. The covariance $\mathbf{\Sigma}$ is a quadratic form on the Lie algebra $\mathfrak{m}$, represented as an 8×8 matrix. As only the $\mathbf{SO}(3)$ and $\mathbf{SO}(2)$ distributions are cyclic, the operator $\rho : \mathfrak{m} \to \mathfrak{m}$, with $\rho((\mathbf{x}, \log r, \mathbf{A}_\mathbf{v}, \mathbf{A}_\theta)) = (\mathbf{0}, 0, \mathbf{A}_\mathbf{v}, \mathbf{A}_\theta)$, causes wrapping to occur only in the rotation components.

Having defined the Gaussian distribution on a single medial atom, the Gaussian distribution of a figure having n medial atoms is the n-fold product distribution on M^n, defined by

$$p(x) = \frac{1}{\sqrt{(2\pi)^{8n}|\mathbf{\Sigma}|}}\sum_{k=-\infty}^{\infty}\exp\Big(-\frac{1}{2}(\mathbf{u}(x)-2\pi k\rho(\bar{\mathbf{u}}(x)))^T\mathbf{\Sigma}^{-1}(\mathbf{u}(x)-2\pi k\rho(\bar{\mathbf{u}}(x)))\Big).$$

Now the vectors $\mathbf{u}(x) = \log(\mu^{-1}x) \in \mathfrak{m}^n$ are $8n$-dimensional, i.e., they are the concatenation of n vectors in $\mathfrak{m}$, representing n medial atoms. The covariance $\mathbf{\Sigma}$ is a quadratic form on $\mathfrak{m}^n$, represented as an $8n \times 8n$ matrix, and the operator ρ projects onto each of the n copies of the rotation groups.

The maximum likelihood estimates for the combined product distribution follow from our development of the maximum likelihood estimates for the individual factors. Recall that the MLE for the mean for each factor is the point that minimizes the sum-of-squared geodesic distances to the sample points. Therefore, the MLE of the mean for samples in the product space is also the minimizer of sum-of-squared geodesic distances. The MLE of the covariance, with the discussed approximations in the rotation dimensions, (10), (15), is the sample covariance matrix in the Lie algebra $\mathfrak{m}^n$. Using an extension of the algorithm in [17] for the intrinsic mean on $\mathbf{SO}(3)$, the intrinsic mean of a collection of m-rep figures with n atoms, $\mathbf{M}_1, \ldots, \mathbf{M}_N \in M^n$, is computed by

Algorithm: M-rep Mean
Input: $\mathbf{M}_1, \ldots, \mathbf{M}_N \in M^n$
Output: $\mu \in M^n$, the mean m-rep
  $\mu = \mathbf{M}_1$
  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 $|| \log(\Delta \mu) || > \epsilon$.

3 Principal Geodesic Analysis

Principal component analysis in $\mathbb{R}^n$ is a powerful technique for analyzing population variation. Principal components of Gaussian data in $\mathbb{R}^n$ are defined as the projection onto the linear subspace spanned by the eigenvectors of the covariance matrix. If we consider a general manifold, the counterpart of a line is a geodesic curve, that is, a curve which minimizes length between two points. In the Lie group M^n geodesics can be computed via the exponential map. Given a tangent vector $\mathbf{v}$ in the Lie algebra $\mathfrak{m}^n$, the geodesic starting at the identity, with initial velocity $\mathbf{v}$, is given by $\gamma : \mathbb{R} \to M^n$, where $\gamma(t) = \exp(t\mathbf{v})$. Similarly, the curve $x \cdot \gamma(t) = x \cdot \exp(t\mathbf{v})$ is a geodesic starting at the point $x \in M^n$.

Since the covariance matrix $\mathbf{\Sigma}$ is a quadratic form on $\mathfrak{m}^n$, its eigenvectors are vectors in the Lie algebra $\mathfrak{m}^n$. These eigenvectors correspond via the exponential map to geodesics on M^n, called *principal geodesics*. The principal geodesic analysis (PGA) on a population of m-rep figures, $\mathbf{M}_i, \ldots, \mathbf{M}_N \in M^n$, is computed by an eigenanalysis of the MLE of the covariance developed above. Thus we have

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

Analogous to linear PCA models, we may choose a subset of the principal directions $\mathbf{u}^{(k)} \in \mathfrak{m}^n$ 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 $[-3\sqrt{\lambda_k}, 3\sqrt{\lambda_k}]$.

4 M-rep Alignment

When applying the above theory for computing means and covariances of real anatomical objects, it is necessary to first globally align the shapes to a common position, orientation, and scale. For objects described by boundary points, the standard method for alignment is the Procrustes method [18]. Procrustes alignment minimizes the sum-of-squared distances between corresponding points. We now develop an analogous alignment procedure based on minimizing sum-of-squared geodesic distances on Lie groups.

Let $\mathbf{S} = (s, \mathbf{R}, \mathbf{w})$ denote a similarity transformation in $\mathbb{R}^3$ consisting of a scaling by $s \in \mathbb{R}^+$, a rotation by $\mathbf{R} \in \mathbf{SO}(3)$, and a translation by $\mathbf{w} \in \mathbb{R}^3$. We define the action of $\mathbf{S}$ on a medial atom $\mathbf{m} = (\mathbf{x}, r, \mathbf{F}, \theta)$ by

$$\mathbf{S} \cdot \mathbf{m} = \mathbf{S} \cdot (\mathbf{x}, r, \mathbf{F}, \theta) = (s\mathbf{R} \cdot \mathbf{x} + \mathbf{w}, sr, \mathbf{R}\mathbf{F}, \theta). \tag{16}$$

Now the action of $\mathbf{S}$ on an m-rep object $\mathbf{M} = \{\mathbf{m}_i : i = 1, \ldots, n\}$ is simply the application of $\mathbf{S}$ to each of $\mathbf{M}$'s medial atoms:

$$\mathbf{S} \cdot \mathbf{M} = \{\mathbf{S} \cdot \mathbf{m}_i : i = 1, \ldots, n\}. \tag{17}$$

It is easy to check from (1) that this action of $\mathbf{S}$ on $\mathbf{M}$ also transforms the implied boundary points of $\mathbf{M}$ by the similarity transformation $\mathbf{S}$.

Consider a collection $\mathbf{M}_1, \ldots, \mathbf{M}_N \in M^n$ of m-rep objects to be aligned, each consisting of n medial atoms. We write $\mathbf{m}_{ij} = (\mathbf{x}_{ij}, r_{ij}, \mathbf{F}_{ij}, \theta_{ij})$ to denote the jth medial atom in the ith m-rep object. Notice that the m-rep parameters, which are positions, rotations, and scalings, are in different units. Before we apply PGA to the m-reps, it is necessary to make the various parameters commensurate. This is done in the Lie algebra by scaling the log rotations and log radii by the average radius value of the corresponding medial atoms. The squared-distance metric between two m-rep models $\mathbf{M}_i$ and $\mathbf{M}_j$ becomes

$$d(\mathbf{M}_i, \mathbf{M}_j)^2 = \sum_{k=1}^{n} (|\mathbf{x}_{jk} - \mathbf{x}_{ik}|^2 + \bar{r}_k^2 (\log r_{jk} - \log r_{ik})^2 + \bar{r}_k^2 |\log(\mathbf{F}_{ik}^{-1} \mathbf{F}_{jk})|^2), \tag{18}$$

where $\bar{r}_k$ is the radius of the kth atom in the mean m-rep. Notice in (16) that the object angle θ is unchanged by a similarity transformation. Thus, the object angles do not appear in the distance metric (18).

The m-rep alignment algorithm finds the set of similarity transforms $\mathbf{S}_1, \ldots, \mathbf{S}_N$ that minimize the total sum-of-squared distances between the m-rep figures:

$$d(\mathbf{S}_1, \ldots, \mathbf{S}_N; \mathbf{M}_1, \ldots, \mathbf{M}_N) = \sum_{i=1}^{N} \sum_{j=1}^{i} d(\mathbf{S}_i \cdot \mathbf{M}_i, \mathbf{S}_j \cdot \mathbf{M}_j)^2. \tag{19}$$

As in generalized Procrustes analysis in $\mathbb{R}^3$, minimization of (19) proceeds in stages:

Algorithm: M-rep Alignment
1. *Translations.* First, the translational part of each $\mathbf{S}_i$ in (19) is minimized once and for all by centering each m-rep model. That is, each model is translated so that the average of it's medial atoms' positions is the origin.
2. *Rotations and Scaling.* The ith model, $\mathbf{M}_i$, is aligned to the mean of the remaining models, denoted μ_i. The alignment is accomplished by a gradient descent algorithm on $\mathbf{SO}(3) \times \mathbb{R}^+$ to minimize $d(\mu_i, \mathbf{S}_i \cdot \mathbf{M}_i)^2$. This is done for each of the N models.
3. *Iterate.* Step 2 is repeated until the metric (19) cannot be further minimized.

5 Results

In this section we present the results of applying our PGA method to a population of 86 m-rep models of the hippocampus from a schizophrenia study. The 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. The models were fit to expert segmentations of the hippocampi from MRI data. The sampling on each m-rep was 3×8, making each model a point on the Lie group M^{24}.

First, the m-rep figures were aligned by the algorithm in §4. The overlayed medial atom centers of the resulting aligned m-reps are shown in Fig. 3. Next, the intrinsic mean m-rep hippocampus was computed (Fig. 3). Finally, PGA was performed on the m-rep figures. The first three modes of variation are shown in Fig. 3.

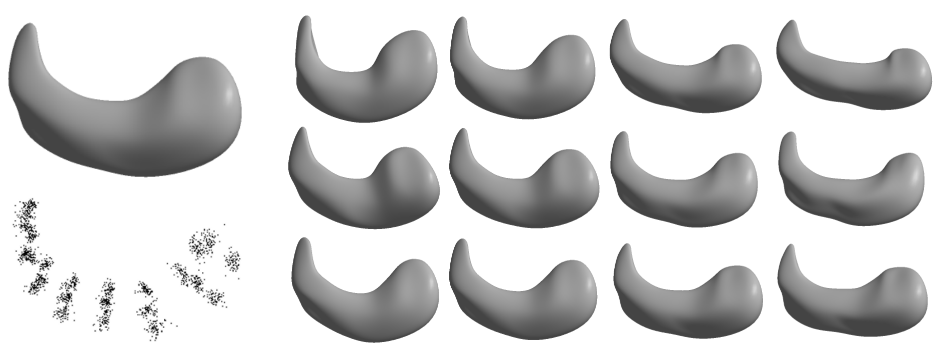

Fig. 3. The surface of the mean hippocampus m-rep (top left). The 86 aligned hippocampus m-reps, shown as overlayed medial atom centers (bottom left). The first three PGA modes of variation for the hippocampus m-reps (right). From left to right are the PGA deformations for -3, -1.5, 1.5, and 3 times $\sqrt{\lambda_i}$.

6 Conclusions

We present a new approach to describing shape variability called principal geodesic analysis. This approach is based on the maximum likelihood estimates of mean and covariance for Gaussian distributions on Lie groups. 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 models that include nonlinear information, such as rotations and magnifications.

7 Acknowledgements

We acknowledge Jeffrey Lieberman and Guido Gerig, UNC Psychiatry, for providing hippocampal data from a clinical Schizophrenia study which was supported by Stanley Foundation and by the UNC-MHNCRC (MH33127). We thank Jeff Townsend for providing the m-rep segmentations of the data. We would like to thank Ulf Grenander, Anuj Srivastava, and Amirjit Budhiraja for useful discussions regarding statistics on Lie groups. This work was done with support from NCI grant P01 CA47982.

References

[1] Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. Computer Vision and Image Understanding **61** (1995) 38–59
[2] Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. In: Fifth European Conference on Computer Vision. (1998) 484–498
[3] Kelemen, A., Szekely, G., Gerig, G.: Three-dimensional model-based segmentation. Transactions on Medical Imaging **18** (1999) 828–839
[4] Jolliffe, I.T.: Principal Component Analysis. Springer-Verlag (1986)
[5] Csernansky, J., Joshi, S., Wang, L., Haller, J., Gado, M., Miller, J., Grenander, U., Miller, M.: Hippocampal morphometry in schizophrenia via high dimensional brain mapping. In: Proceedings National Academy of Sciences. (1998) 11406–11411
[6] Joshi, S., Pizer, S., Fletcher, P.T., Yushkevich, P., Thall, A., Marron, J.S.: Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. Transactions on Medical Imaging **21** (2002)
[7] Blum, H., Nagel, R.: Shape description using weighted symmetric axis features. Pattern Recognition **10** (1978) 167–180
[8] Thall, A.: Fast C^2 interpolating subdivision surfaces using iterative inversion of stationary subdivision rules. (2002) http://midag.cs.unc.edu/pub/papers/Thall_TR02-001.pdf.
[9] Duistermaat, J.J., Kolk, J.A.C.: Lie Groups. Springer (2000)
[10] Curtis, M.L.: Matrix Groups. Springer-Verlag (1984)
[11] Murray, R., Li, Z., Sastry, S.: A Mathematical Introduction to Robotic Manipulation. CRC Press (1994)
[12] Mardia, K.V.: Directional Statistics. John Wiley and Sons (1999)
[13] Kendall, D.G.: Shape manifolds, Procrustean metrics, and complex projective spaces. Bulletin of the London Mathematical Society **16** (1984) 18–121
[14] Moakher, M.: Means and averaging in the group of rotations. SIAM Journal on Matrix Analysis and Applications **24** (2002) 1–16
[15] Grenander, U.: Probabilities on Algebraic Structures. John Wiley and Sons (1963)
[16] Lee, J.M.: Riemannian Manifolds: An Introduction to Curvature. Springer-Verlag (1997)
[17] Buss, S.R., Fillmore, J.P.: Spherical averages and applications to spherical splines and interpolation. ACM Transactions on Graphics **20** (2001) 95–126
[18] Goodall, C.: Procrustes methods in the statistical analysis of shape. Journal of the Royal Statistical Society **53** (1991) 285–339
[19] Styner, M., Gerig, G.: Medial models incorporating object variability for 3D shape analysis. In: Information Processing in Medical Imaging. (2001) 502–516