Published in final edited form as:

Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit. 2012 June ; 2012: 1027–1034. doi:10.1109/CVPR.2012.6247780.

Sasaki Metrics for Analysis of Longitudinal Data on Manifolds

Prasanna Muralidharan and P. Thomas Fletcher

School of Computing, University of Utah

Prasanna Muralidharan: prasanna@sci.utah.edu; P. Thomas Fletcher: fletcher@sci.utah.edu

Abstract

Longitudinal data arises in many applications in which the goal is to understand changes in individual entities over time. In this paper, we present a method for analyzing longitudinal data that take values in a Riemannian manifold. A driving application is to characterize anatomical shape changes and to distinguish between trends in anatomy that are healthy versus those that are due to disease. We present a generative hierarchical model in which each individual is modeled by a geodesic trend, which in turn is considered as a perturbation of the mean geodesic trend for the population. Each geodesic in the model can be uniquely parameterized by a starting point and velocity, i.e., a point in the tangent bundle. Comparison between these parameters is achieved through the Sasaki metric, which provides a natural distance metric on the tangent bundle. We develop a statistical hypothesis test for differences between two groups of longitudinal data by generalizing the Hotelling T^2 statistic to manifolds. We demonstrate the ability of these methods to distinguish differences in shape changes in a comparison of longitudinal corpus callosum data in subjects with dementia versus healthily aging controls.

1. Introduction

A longitudinal study tracks changes in individuals by repeatedly collecting measurements over time. Longitudinal studies are popular in medicine, where the goal is to understand change processes, such as healthy development, aging, or disease progression. Often, shape is the quantity of interest being tracked. For example, understanding changes in neuroanatomy is a critical goal in the study of degenerative diseases such as Alzheimer's and in developmental disorders such as autism. Longitudinal shape data also arises in various branches of biology, such as evolutionary biology, where the evolution of the shapes of bones in the fossil record is of interest. The main challenge for these studies is that shape, i.e., the geometry of an object that is invariant to rotation, scaling, and translation, is inherently nonlinear and high-dimensional. Because of this, manifold representations of shape have proven to be effective. Therefore, analysis of shape changes necessitates the development of models for dealing with manifold-valued longitudinal data. Such models would also benefit other applications that involve serial collection of manifold data, such as directional data, transformation groups, and tensors.

Related to the longitudinal data analysis problem is the regression problem. However, regression does not model *individual* changes and is not appropriate for analyzing longitudinal data. Instead, regression models are used for describing cross-sectional data, where only one data point per individual is available. Several authors have proposed

methods for regression on manifolds. Jupp and Kent [6] propose an unrolling method on shape spaces. Regression analysis on the group of diffeomorphisms has been proposed as growth models by Miller [10], nonparametric regression by Davis, et al. [1], and second-order splines by Trouvé and Vialard [16]. Shi, et al. [15] proposed a semiparametric model with multiple covariates for manifold response data. Recently, parametric models of regression, where the regression function is a geodesic curve, have been introduced independently by Fletcher [5] and Niethammer et al. [11].

Related work in longitudinal analysis includes several approaches in the setting of diffeomorphic transformations, which form an infinite-dimensional manifold, applied to image sequences. Durrleman et al. [4] construct spatiotemporal image atlases from longitudinal data. Qiu et al. [13] use parallel translation to bring individual trajectories to a common point for comparison. Lorenzi et al. [9] use a hierarchical model on stationary velocity fields, in a framework that does not include a Riemannian metric on the manifold of diffeomorphisms. An important shortcoming of these approaches is that they do not model *distances* between trajectories. This makes it difficult to compare the differences in trends of two groups, or even to rigorously define the concept of the variance of a population of trends.

We propose a generative hierarchical model for longitudinal data analysis on Riemannian manifolds. Serial data from each individual is represented by a geodesic trend in the first stage, and these trends are in turn modeled as perturbations of the mean geodesic trend of the group in the second stage. Geodesic trends in both stages are uniquely parameterized by their initial conditions: an initial position, or “intercept”, and an initial velocity, or “slope”. We then define a distance metric between trends, which allows us to develop least-squares estimation of our model parameters, a definition of the variance of trends, and a method for comparing the mean trends between two groups. We do this by considering the slope-intercept pairs to be elements of the tangent bundle, which can be given the structure of a differentiable manifold with geodesic distances defined by the Sasaki metric.

2. Riemannian Geometry Preliminaries

Before introducing our longitudinal model, we briefly review some necessary facts about Riemannian geometry (see [2] for more details). Recall that a Riemannian manifold (M, g) is a differentiable manifold M equipped with a metric g , which provides a smoothly varying inner product on the tangent spaces of M . Given two vector fields u, w on M , the covariant derivative $\nabla_u w$ gives the change of the vector field w in the u direction. The covariant derivative is a generalization of the Euclidean directional derivative to the manifold setting. Consider a curve $\gamma: [0, 1] \rightarrow M$ and let $\dot{\gamma} = d\gamma/dt$ be its velocity. Given a vector field $V(t)$

defined along γ , we can define the covariant derivative of V to be $\frac{DV}{dt} = \nabla_{\dot{\gamma}} V$. A vector field is called *parallel* if the covariant derivative along the curve γ is zero. A curve γ is geodesic if it satisfies the equation $\nabla_{\dot{\gamma}} \dot{\gamma} = 0$. In other words, geodesics are curves with zero acceleration.

We write an element of the tangent bundle as the pair $(p, u) \in TM$, where p is a point in M and $u \in T_p M$ is a tangent vector at p . The tangent bundle TM can also be given the structure of a differentiable manifold, which is twice the dimension of the original manifold M . The tangent bundle TM serves as a convenient parametrization of the set of possible geodesics on M . Recall that for any $(p, u) \in TM$ there is a unique geodesic curve γ , with initial conditions $\gamma(0) = p$ and $\dot{\gamma}(0) = u$. This geodesic is only guaranteed to exist locally. When γ is defined over the interval $[0, 1]$, the exponential map at p is defined as $\text{Exp}_p(u) = \gamma(1)$. In other words, the exponential map takes a position and velocity as input and returns the point at time 1 along the geodesic with these initial conditions. The exponential map is locally diffeomorphic onto a neighborhood of p . Let $V(p)$ be the largest such neighborhood. Then within $V(p)$ the exponential map has an inverse, the Riemannian log map, $\text{Log}_p : V(p) \rightarrow T_p M$. For any point $q \in V(p)$, the Riemannian distance function is given by $d(p, q) = \|\text{Log}_p(q)\|$. It will be convenient to include the point p as a parameter in the exponential and log maps, i.e., define $\text{Exp}(p, u) = \text{Exp}_p(u)$ and $\text{Log}(p, q) = \text{Log}_p(q)$.

If u, v, w are vector fields on a Riemannian manifold M , the Riemannian curvature tensor R on M is defined as

$$R(u, v)w = \nabla_v \nabla_u w - \nabla_u \nabla_v w + \nabla_{[u, v]} w$$

where $[u, v]$ is the Lie bracket of the pair of vector fields u and v . Intuitively, for a manifold M , the curvature R measures how far away it is from being “flat”, i.e., the manifold being locally isometric to the Euclidean space. The curvature tensor for the Euclidean space evaluates to zero, and for a general manifold, R measures the extent of non-commutativity of the covariant derivative.

3. Longitudinal Analysis on Manifolds

Our model is inspired by the work of Laird and Ware [8], who proposed a hierarchical mixed-effects model for longitudinal data in a linear vector space. Consider a longitudinal response variable for the i th individual, Y_i , taking values in a Riemannian manifold M . Realizations of the response will be denoted y_{ij} corresponding to the j th observation of the i th individual. Let X_i denote the independent variable, typically time, with realizations $x_{ij} \in \mathbb{R}$, corresponding to each y_{ij} . We propose the following hierarchical model for manifold-valued longitudinal data (see Figure 1).

Individual Level—Each individual response, Y_i , is modeled by a geodesic trend with noise (cf. [5]),

$$Y_i = \text{Exp}(\text{Exp}(p_i, X_i u_i), \varepsilon_i). \quad (1)$$

The inner exponential map is a geodesic curve for the i th individual, determined by an initial position, p_i , and velocity, u_i , and parameterized with respect to the independent variable, X_i . The outer exponential map models the random variation of observations from this curve,

where ε_i is a random variable taking values in the tangent space at the corresponding position on the geodesic, $\text{Exp}(p_i, X_i u_i)$.

Group Level—The parameters, (p_i, u_i) , of the individual trends are modeled as perturbations from a mean geodesic trend for the population, which is determined by the parameters $(\alpha, \beta) \in TM$. The group model is given by

$$(p_i, u_i) = \text{Exp}_S((\alpha, \beta), (v_i, w_i)), \quad (2)$$

where (v_i, w_i) is a tangent vector of TM , representing the change in the α and β parameters, and Exp_S is the exponential map on the tangent bundle TM , equipped with the Sasaki metric described below.

Figure 1 gives a diagram of the proposed model. The observations y_{ij} for different individuals are drawn with different symbols. The dashed red curves depict the individual geodesic trends given by (1), and the blue curve represents the mean trend of the population given by (2).

To define the random variation of the individual trends in the group level model (2) above, we equip the tangent bundle TM with a Riemannian metric, called the Sasaki metric [14]. This makes possible the geodesic perturbations from the mean trend, given by the exponential map Exp_S on TM . Let the pair (v, w) represent a tangent vector to TM at a point (p, u) . Intuitively, v represents a change in p , and w represents a change in u . Technically, v and w live in the tangent space $T_p M$, and they need to be “lifted” to $T_{(p, u)} TM$, i.e., the tangent space to TM at (p, u) . The lift of the v component, denoted v^h , is called the *horizontal lift* of v . Geodesics along v^h changes the point p while parallel translating u . The lift of w , denoted w^v , is called the *vertical lift* of w . Geodesics along w^v leave p fixed and move u linearly. We will use the notation $(v, w) \equiv v^h + w^v$.

Now, to define a metric on the tangent bundle TM , we need to define the inner product between combinations of horizontal and vertical components. Given two tangent vectors $a = (v_1, w_1), b = (v_2, w_2) \in T_{(p, u)} TM$, the Sasaki metric $\bar{g}(a, b)$, is given as

$$\begin{aligned} \bar{g}(v_1^h, v_2^h) &= g(v_1, v_2), \\ \bar{g}(v_1^h, w_2^v) &= \bar{g}(v_2^h, w_1^v) = 0, \\ \bar{g}(w_1^v, w_2^v) &= g(w_1, w_2), \end{aligned}$$

where g is the metric on M .

Let $\eta(t) = (p(t), u(t))$ be a geodesic curve in TM . Then η satisfies the geodesic equation

$$\bar{\nabla}_{\dot{\eta}} \dot{\eta} = 0,$$

where $\bar{\nabla}$ is the covariant derivative defined by the Sasaki metric on TM . Splitting $\dot{\eta}$ into its horizontal and vertical components, $\dot{\eta}(t) = v^h(t) + w^v(t)$, the geodesic equation can be written as a pair of coupled equations,

$$\nabla_v v = -R(u, w)v, \quad (3)$$

$$\nabla_v w = 0, \quad (4)$$

where R is the Riemannian curvature tensor of the manifold M . Notice that (3) says that the position component, the $p(t) \in M$, will bend depending on the curvature tensor and the vectors u, v, w , while (4) indicates that the $u(t) \in T_{p(t)}M$ component will change at a constant rate given by the parallel vector field $w(t)$.

3.1. Least-Squares Estimation

To estimate the parameters of the geodesic longitudinal model in (1), (2), we use a two-step least-squares estimation procedure. The least-squares problem is phrased in terms of the sum-of-squared Riemannian distance between the data and the model.

The first step is to estimate the slope and intercept $(p_i, u_i) \in TM$ for each individual, given data $(x_{ij}, y_{ij}) \in \mathbb{R} \times M$. Following the geodesic regression models in [5, 11] we minimize the sum-of-squared residuals, giving

$$(\hat{p}_i, \hat{u}_i) = \arg \min_{(p, u)} \sum_{j=1}^{n_i} d(\text{Exp}(p, x_{ij}u), y_{ij})^2. \quad (5)$$

This minimization problem can be solved using gradient descent, where the gradients are given in terms of Jacobi fields along the current estimate of the geodesic. See [5] for more details.

In the second step, we estimate the parameters (α, β) in the group model (2). Again we take a least-squares approach, where now we want to minimize the sum-of-squared geodesic distances between the estimates $(\hat{\alpha}, \hat{\beta})$ and the individual trends (p_i, u_i) . Now geodesic distance is in the tangent bundle TM , with respect to the Sasaki metric. This amounts to computing the Fréchet expectation of the (p_i, u_i) in TM , which is defined as

$$(\hat{\alpha}, \hat{\beta}) = \arg \min_{(p, u)} \sum_{i=1}^N d_S((p_i, u_i), (p, u))^2, \quad (6)$$

where d_S denotes geodesic distance in TM under the Sasaki metric. A gradient descent procedure for minimizing this equation is as follows:

Algorithm 1: Group Parameter Estimation

Initialize: $(\hat{\alpha}, \hat{\beta})_0 = (p_1, u_1)$

Iterate over k :

$$\begin{aligned}\delta &= \frac{\tau}{N} \sum_{i=1}^N \text{Log}_S((\hat{\alpha}, \hat{\beta})_k, (p_i, u_i)) \\ (\hat{\alpha}, \hat{\beta})_{k+1} &= \text{Exp}_S((\hat{\alpha}, \hat{\beta})_k, \tau \delta)\end{aligned}$$

While $\|\delta\| > \varepsilon$.

The update δ in this algorithm is the negative gradient of (6) at the current estimate, τ is a step size, and ε is the stopping criteria. This algorithm requires that we have a procedure to compute the exponential and log maps on TM under the Sasaki metric, which we now develop.

3.2. Computing Sasaki Geodesics on TM

In this section, we give algorithms for computing the Exp_S and Log_S maps on TM with the Sasaki metric. For most manifolds, these maps will not have closed form. But for many useful manifolds, we do have closed form solutions for computations on M . Thus, we formulate the discretized form of the geodesic equation on TM , in terms of computations involving geodesics, parallel translation, and the curvature tensor on M .

An example of a Sasaki geodesic on TS^2 , the tangent bundle of the sphere, is shown in Figure 2. This was computed using the Log_S map algorithm described below. A geodesic curve $(p(t), u(t))$ on the tangent bundle TS^2 can be thought of as a curve on the sphere $p(t)$, shown in black, and a corresponding vector field $u(t)$ along that curve, shown in red. Notice that a geodesic on S^2 would be a great circle, and that the $p(t)$ curve shown here is not a geodesic on S^2 , but rather must bend according to the curvature tensor term in (3).

Exponential Map—The Sasaki exponential map computes a geodesic on TM with input initial conditions $(p_0, u_0) \in TM$ for the initial position and $(v_0, w_0) \in T_{(p_0, u_0)} TM$ for the initial velocity. We shoot a Sasaki geodesic using an Euler integration on TM . Let $k = 1, \dots, L$ be the discrete time step and $\varepsilon = 1/L$. The update equations for k th step is

$$\begin{aligned}p_{k+1} &= \text{Exp}(p_k, \varepsilon v_k), \\ u_{k+1} &= \phi(u_k + \varepsilon w_k, p_{k+1}), \\ v_{k+1} &= \phi(v_k - \varepsilon R(u_k, w_k)v_k, p_{k+1}), \\ w_{k+1} &= \phi(w_k, p_{k+1}),\end{aligned}$$

where $\phi(v, p)$ denotes parallel translation of the vector v to the point p . Notice that p_{k+1} is simply a small step M -exponential map in the direction of v_k . All other computations are tangent vector computations. Once computed, they need to be parallel translated to the next point, p_{k+1} . Again, the exponential map, parallel translation, and curvature tensor for M are assumed to be known. We give specific formulae for these when we consider example manifolds in the next section.

Log Map—Starting with points $a = (p_0, u_0) \in TM$ and $b = (p_L, u_L) \in TM$, the Sasaki log map, $\text{Log}_S(a, b)$, returns the tangent $(v, w) \in T_a TM$, which is the initial velocity of the geodesic segment between a and b . Now we will iteratively relax a discretized geodesic between a and b , denoted $\{(p_k^i, u_k^i)\}$, where $k = 0, \dots, L$ is the discrete time step along the curve, and i is the current iteration. The relaxation procedure minimizes the discrete geodesic energy,

$$E(p, u) = \sum_{k=1}^L (\|v_k\|^2 + \|w_k\|^2),$$

where v_k is the tangent vector of a geodesic segment on M between p_k and p_{k+1} , and w_k is the linear change between u_k and u_{k+1} after parallel translation.

Minimization of the discrete geodesic energy proceeds by gradient descent as follows:

- The initial curve is obtained by discretizing the M -geodesic between p_0 and p_L , and linearly interpolating between u_0 and u_L in parallel translated coordinates as

$$u_k = (1 - k/L) \phi(u, p_k) + (k/L) \phi(u_L, p_k).$$

- Compute the v_k^i and w_k^i as finite differences:

$$v_k^i = \text{Log}(p_k^i, p_{k+1}^i) / \varepsilon \quad (7)$$

$$w_k^i = (\phi(u_{k+1}^i, p_k^i) - \phi(u_{k-1}^i, p_k^i)) / (2\varepsilon) \quad (8)$$

- Compute $\nabla_{v_k^i} v_k^i$ as a central difference,

$$\nabla_{v_k^i} v_k^i = (\phi(v_{k+1}^i, p_k^i) - \phi(v_{k-1}^i, p_k^i)) / (2\varepsilon)$$

$\nabla_{v_k^i} w_k^i$ is similarly computed as a second-order finite difference of u_k^i 's (i.e., repeated first-order finite differences).

- The gradients of E with respect to p_k^i and u_k^i are

$$\begin{aligned} \nabla_{p_k^i} E &= \nabla_{v_k^i} v_k^i + R(u_k^i, w_k^i) v_k^i \\ \nabla_{u_k^i} E &= \nabla_{v_k^i} w_k^i \end{aligned}$$

To clarify notation, the ∇ on the left is the gradient of the energy, whereas the ∇ on the right denotes the covariant derivative of tangent vectors along other tangent vectors.

- Now update the discrete curve in the negative gradient direction:

$$\begin{aligned} p_k^{i+1} &= \text{Exp}(p_k^i, \delta \nabla_{p_k}^i E) \\ u_k^{i+1} &= \phi(u_k^i + \delta \nabla_{p_k}^i E, p_k^{i+1}) \end{aligned}$$

Given the converged discrete geodesic, the output of the Log_S map is the initial tangent (v, w), given by

$$\begin{aligned} v &= \text{Log}(p_0, p_1)/\varepsilon \\ w &= (\phi(u_1, p_0) - u_0)/\varepsilon \end{aligned}$$

3.3. Testing Group Differences

One of the major motivations of longitudinal data analysis is to test if changes observed in one group differ from those found in another. For instance, one might ask if the brain anatomy of Alzheimer's patients deteriorates faster than those of healthily aging subjects. In this section we develop a statistical hypothesis test for comparing the Sasaki average trends between two groups. We do this by generalizing the Hotelling T^2 statistic to the manifold setting, and applying this to the tangent bundle TM equipped with the Sasaki metric. To test the statistical significance of the group difference, we use a permutation test on this generalized T^2 statistic.

Recall the Hotelling T^2 statistic is a multivariate test of the difference between sample means, $\bar{p}, \bar{q}$, of two groups of data $\{p_1, \dots, p_m\}$ and $\{q_1, \dots, q_n\}$, with all $p_i, q_i \in \mathbb{R}^d$. The idea is to compare the difference between the two means, relative to the pooled sample covariance:

$$W = \frac{\sum_i (p_i - \bar{p})(p_i - \bar{p})^T + \sum_i (q_i - \bar{q})(q_i - \bar{q})^T}{m+n-2}.$$

The T^2 statistic can be thought of as a squared Mahalanobis distance between the means, using this pooled covariance, W . The sample T^2 statistic is given by

$$t^2 = \frac{mn}{m+n} (\bar{p} - \bar{q})^T W^{-1} (\bar{p} - \bar{q})$$

To generalize the Hotelling T^2 statistic to the manifold setting, consider two samples $\{p_1, \dots, p_m\}$ and $\{q_1, \dots, q_n\}$, with all p_i, q_i now being points on some Riemannian manifold N . For the purposes of analyzing longitudinal trends, we will mostly be interested in data on a tangent bundle manifold $N = TM$ equipped with the Sasaki metric. Now we wish to develop a statistic for testing the differences between the sample Fréchet means, $\bar{p}, \bar{q}$, of the two groups. Note that $\bar{p}, \bar{q}$ are computed by solving the minimization problem (6). The difference between the means can be represented as the tangent vector $v_p = \text{Log}(p, \bar{q})$, or as the vector $v_q = \text{Log}(q, \bar{p})$. However, the difficulty is that these two options live in two different tangent spaces, $v_p \in T_p N$ and $v_q \in T_q N$, respectively. Similarly, the sample covariance matrices for

the p_i and q_i are defined in these different tangent spaces. Therefore, pooling the sample covariance matrix is not straightforward. Instead we compute two mean differences, one in each tangent space and weighted by the respective single-group covariance, and then average the results. This gives the following generalization of the sample T^2 statistic:

$$t^2 = \frac{1}{2} (v_p^T W_p^{-1} v_p + v_q^T W_q^{-1} v_q), \quad (9)$$

where the individual group covariances are computed as

$$W_p = \frac{1}{m} \sum_i \text{Log}(\bar{p}, p_i) \text{Log}(\bar{p}, p_i)^T, \\ W_q = \frac{1}{n} \sum_i \text{Log}(\bar{q}, q_i) \text{Log}(\bar{q}, q_i)^T.$$

To test the statistical significance of the manifold T^2 statistic, we use a nonparametric permutation test. The motivation for this is two-fold. First, it is difficult to formulate a parametric distribution for data on a general manifold and then derive the resulting parametric distribution of the T^2 statistic. Second, even if we had such a parametric formulation, we prefer to not make such assumptions about the distribution of the data we are given. The permutation test procedure is as follows: (1) compute the t^2 statistic, (2) randomly permute (swap) data points between the p and q groups, computing a t_k^2 statistic for the permuted groups, (3) repeat step 2 for $k = 1, \dots, P$, (4) compute the p -value: $p = B/(P + 1)$, where B is the number of $t_k^2 < t^2$. The final p -value can be interpreted as the probability of finding a larger group difference by random chance under the null hypothesis (that there is no difference between the means).

We now return to the specific problem of comparing the mean trends in two different groups. Consider two sets of longitudinal data $y_{ij}^{(1)}$ and $y_{ij}^{(2)}$ on M and the resulting parameter estimates for the two groups, $(\hat{\alpha}_1, \hat{\beta}_1)$ and $(\hat{\alpha}_2, \hat{\beta}_2)$, using the hierarchical model and estimation described above. It is often most interesting to separate the tests of the intercept parameter α and the slope parameter β . For example, in testing the differences in anatomical changes between a healthy and disease group, it is important to distinguish if the shape differences are present at baseline (intercept) or if they develop over time (slope). To make this distinction, we can separate the above Hotelling T^2 test into these two components. Let (v, w) be a tangent vector to TM and define the projection operators $\pi_1(v, w) = (v, 0)$ and $\pi_2(v, w) = (0, w)$. The two separated statistics, t_α^2 and t_β^2 , are now given by (9), with the change that all tangent vectors are replaced with their projected versions, using π_1 for the α test and π_2 for the β test. Note that this includes projection of the vectors used in computing the covariances W_p and W_q .

4. Results

We validate the proposed model and estimation procedure on two example manifolds. First, we generated data from the geodesic longitudinal model on the sphere, S^2 , and tested the

least-squares estimation procedure with this generated data. Second, we apply the geodesic longitudinal model and Hotelling T^2 statistic to test group differences in a real longitudinal data set of corpus callosum shapes in individuals with and without dementia.

Review of Sphere Geometry

Let p be a point on an n -dimensional sphere embedded in $\mathbb{R}^{n+1}$, and let u be a tangent at p . Let the inner product defined between tangents at a base point p , be the usual Euclidean inner product. The exponential map is then given by a 2D rotation on the sphere by an angle given by the norm of the tangent, given as,

$$\text{Exp}_p(u) = \cos \theta \cdot p + \frac{\sin \theta}{\theta} \cdot u, \quad \theta = \|u\|. \quad (10)$$

Likewise, the log map between two points p and q on the sphere can be done by finding the initial velocity of the rotation between the two points. Let $\pi_p(q) = p \cdot \langle p, q \rangle$ denote the projection of the vector q onto p . Then the log map is given by

$$\text{Log}_p(q) = \frac{\theta \cdot (q - \pi_p(q))}{\|q - \pi_p(q)\|}, \quad \theta = \arccos(\langle p, q \rangle). \quad (11)$$

Let u, v, w be tangent vectors at a point p , on the sphere. Then the Riemannian curvature tensor for the sphere can be written as

$$R(u, v)w = \langle w, u \rangle v - \langle w, v \rangle u \quad (12)$$

Synthetic Longitudinal Sphere Experiment

To test the estimation procedure in Section 3.1, we generated synthetic data on S^2 from our hierarchical longitudinal model given in (1) and (2). We started with fixed parameters $\alpha =$

$(1, 0, 0)$ $\beta = (0, \frac{\pi}{4}, 0)$ as the group intercept and slope. We then generate (p_i, u_i) via the group model given by (2). In this case, we generated 30 of these points on TS^2 via the Sasaki exponential map for each i , where the (v_i, w_i) , i.e., the vector perturbations about the mean

trend, were taken to be isotropic Gaussian with mean zero and $\sigma = \frac{\pi}{16}$. Once we had (p_i, u_i) , we generated 8 time points for each individual, with x_{ij} data from a uniform distribution on

$[0, 1]$, and isotropic Gaussian tangent vector residuals, ε_{ij} , with mean zero and $\sigma = \frac{\pi}{16}$. We then used the sphere exponential map to generate the data points y_{ij} via the model (1).

Having generated data, we ran the least-squares estimation procedure described in Section 3.1 to test whether we can faithfully recover the same parameters α, β from which the data were generated. The results showed that the distance between the estimated $(\hat{\alpha}, \hat{\beta})$ and the true (α, β) was 0.035. The resulting estimated model is shown in Figure 3. The red vectors are the estimated (p_i, u_i) for each individual, the blue vector is the estimated group

parameters, $(\hat{\alpha}, \hat{\beta})$, and the black vector is the true values for (α, β) from which the model was generated.

Review of Kendall's Shape Space

Intuitively, shape can be thought of as the geometric properties of an object that remain, when location, scale and rotational effects are removed. We use the shape space introduced by Kendall [7], which is the Riemannian manifold formed by taking the equivalence classes of configurations of k points in $\mathbb{R}^2$ under translation, rotation, and scaling.

A specific case would be a configuration of k points in the 2D plane. This can be represented as a complex k -vector, $z \in \mathbb{C}^k$. Removing translation means requiring the centroid to be zero, which projects this point to the linear complex subspace $V = \{z \in \mathbb{C}^k : \sum z_i = 0\}$, which is equivalent to the space $\mathbb{C}^{k-1}$. Next, points in this subspace are deemed equivalent if they are a rotated and a scaled version of each other. This can be represented as multiplication by a complex number, $\rho e^{i\theta}$, where ρ is the scaling factor and θ is the rotation angle. The set of such equivalence classes forms the 2D shape space, also known in our representation as the complex projective space, $\mathbb{C}P^{k-2}$.

Thus, we think of a centered shape $p \in V$ as representing the complex line $L_p = \{z \cdot p : z \in \mathbb{C} \setminus \{0\}\}$, i.e., L_p consists of all point configurations with the same shape as p . A tangent vector at $L_p \in V$ is a complex vector, $u \in V$, such that $\langle p, u \rangle = 0$.

The nice thing about the Kendall shape space is that it forms a Riemannian manifold. The exponential map is given by rotating (within V) the complex line L_p by the initial velocity u , that is,

$$Exp_p(u) = \cos \theta \cdot p + \frac{\|p\| \sin \theta}{\theta} \cdot u, \quad \theta = \|u\|. \quad (13)$$

Likewise, the log map between two shapes $p, q \in V$ is given by finding the initial velocity of the rotation between the two complex lines L_p and L_q . We first Procrustes align q to p by computing $z = pq$, which is the rotation needed for this alignment. Let $\pi_p(zq) = p \cdot \langle p, zq \rangle / \|p\|^2$ denote the projection of the vector zq onto p . Then the log map is given by

$$Log_p(q) = \frac{\theta \cdot (zq - \pi_p(zq))}{\|zq - \pi_p(zq)\|}, \quad \theta = \arccos \frac{|\langle p, zq \rangle|}{\|p\| \|zq\|}. \quad (14)$$

Notice that we never explicitly project a shape onto $\mathbb{C}P^{k-2}$. This has the effect that shapes computed via the exponential map (13) will have the same orientation and scale as the base point p . Also, tangent vectors computed via the log map (14) are valid only at the particular representation p (and not at a rotated or scaled version of p). This works nicely for our purposes and implies that shapes along the estimated geodesic will have the same orientation and scale as the intercept shape, $\hat{p}$.

The complex projective space is a Riemannian manifold with non-constant curvature. The Riemannian curvature tensor of $\mathbb{C}P^{k-2}$ can be computed as follows. Let u, v, w be vectors at a point $p \in \mathbb{C}P^{k-2}$. These vectors can be represented in $\mathbb{C}^{k-1} \cong \mathbb{R}^{2k-2}$. Writing the vector w as $w = (w_1, \dots, w_{2k-2})$, define the operator

$$Jw = (-w_k, \dots, -w_{2k-2}, w_1, \dots, w_{k-1}).$$

(This is just multiplication by $i = \sqrt{-1}$ if we take w as a complex vector with the $k-1$ real coordinates listed first.)

Using this operator, the curvature tensor R^* is given as

$$\begin{aligned} R^*(u, v)w &= R(u, v)w + \langle v, Jw \rangle Ju \\ &+ \langle w, Ju \rangle Jv - 2\langle u, Jv \rangle Jw \end{aligned}$$

where R is the curvature tensor of the sphere, S^{2k-1} . For more details, refer to [12].

Corpus Callosum Shape Changes in Healthy Aging vs. Dementia

The corpus callosum is the major white matter bundle connecting the two hemispheres of the brain. A midsagittal slice from a magnetic resonance image (MRI) with segmented corpus callosum is shown in Figure 4. Several studies have shown that the volume of the corpus callosum decreases with normal aging [3]. There have also been studies involving regional measurements, such as volume, length, and local curvature. However, there hasn't been a longitudinal study of individuals with and without dementia that takes into account the *entire* shape of the corpus callosum. The advantage of analyzing anatomical data in a shape space is that it takes into account all shape properties, and their correlations, at once. There is no need to predefine derived shape measurements (length, curvature, etc.).

The longitudinal data we used was from the OASIS brain database (<http://www.oasis-brains.org/>). This database has about 150 subjects, aged between 60 to 96 years old, with MRI for two or more time points. Each time point is separated by at least one year. Each subject is characterized as non-demented, demented, or as having converted from non-demented to demented during the study. For our study, we worked with only the male subjects, to avoid gender effects. Also, we only worked with subjects having MRI for at least 3 different time-points. These selection criteria resulted in 69 total corpus callosum shapes from 11 subjects with dementia and 12 without dementia, each with 3 time points.

Figure 5 shows the result of the geodesic longitudinal model estimation for the corpus callosum shape trends in the dementia (top) and non-dementia (bottom) groups. The estimated model parameters can be interpreted as the average shape at baseline (the intercept, $\hat{\alpha}$) and the change in shape over the 6 year course of the study (the slope, $\hat{\beta}$). The average shape at the end six year time period (shown in red for each group) was generated by shooting along a geodesic starting at $\hat{\alpha}$ with initial velocity $\hat{\beta}$. Note that the average baseline shapes for the two groups are quite similar. However, there is a stark difference in

the changes in the corpus callosum over time, with the dementia group displaying more drastic thinning and bending over six years.

To test the statistical significance of the difference between the two groups, we performed a permutation test with the manifold Hotelling T^2 test as described in Section 3.3. Differences in the mean intercepts, $\hat{\alpha}$, and the mean slopes, $\hat{\beta}$, were tested separately. We computed t^2 statistics for 10,000 permutations and computed p -values for both intercept and slope differences. These results are shown in Table 1. The difference in the mean intercept parameters between the two groups was not found to be significant. This matches the fact that the baseline shapes in Figure 5 look fairly similar. However, comparing the two groups on the basis of the slope parameter, i.e, how the corpus callosum shape has changed over time, we found a significant difference. Again, this matches the obvious differences seen in the end point shapes seen in Figure 5.

5. Conclusion

We proposed a generative, hierarchical model for longitudinal data analysis of serial manifold-valued data. A key property of this model is the power of comparing geodesic trends on the manifold using the Sasaki metric of the tangent bundle. This provides methods for estimating average geodesic trends and for comparing differences between two groups of longitudinal manifold data. The power of the model to distinguish group differences in shape trends has been demonstrated in a comparison of longitudinal corpus callosum data of subjects with and without dementia. While the geodesic model proved effective for this application, future work will include development of more flexible curve models for individual trends on manifolds.

Acknowledgments

This work was supported by NSF CAREER Award 1054057.

References

1. Davis, B.; Fletcher, P.T.; Bullitt, E.; Joshi, S. Population shape regression from random design data. *Proceedings of IEEE International Conference on Computer Vision*; 2007.
2. do Carmo, M. *Riemannian geometry*. Birkhäuser; 1992.
3. Driesen N, Raz N. The influence of sex, age, and handedness on corpus callosum morphology: A meta-analysis. *Psychobiology*. 1995; 23(3):240–247.
4. Durrleman S, Pennec X, Trounev A, Gerig G, Ayache N. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. *Medical Image Computing and Computer-Assisted Intervention*. 2009:297–304. [PubMed: 20426000]
5. Fletcher P.T. Geodesic regression on Riemannian manifolds. *MICCAI Workshop on Mathematical Foundations of Computational Anatomy*. 2011:75–86.
6. Jupp P.E., Kent J.T. Fitting smooth paths to spherical data. *Applied Statistics*. 1987; 36(1):34–46.
7. Kendall D.G. Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*. 1984; 16:18–121.
8. Laird N, Ware J.H. Random-effects models for longitudinal data. *Biometrics*. 1982; 38(4):963–974. [PubMed: 7168798]

9. Lorenzi M, Ayache N, Frisoni GB, Pennec X. Mapping the effects of Ab142 levels on the longitudinal changes in healthy aging: hierarchical modeling based on stationary velocity fields. *Proceedings of Medical Image Computing and Computer Assisted Intervention*. 2011
10. Miller M. Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *NeuroImage*. 2004; 23:S19–S33. [PubMed: 15501089]
11. Niethammer M, Huang Y, Vialard F-X. Geodesic regression for image time-series. *Medical Image Computing and Computer Assisted Intervention*. 2011
12. O’Neill B. The fundamental equations of a submersion. *Michigan Math J*. 1966; 13:459–469.
13. Qiu A, Albert M, Younes L, Miller M. Time sequence diffeomorphic metric mapping and parallel transport track time-dependent shape changes. *NeuroImage*. 2009; 45:S51–S60. [PubMed: 19041947]
14. Sasaki S. On the differential geometry of tangent bundles of riemannian manifolds. *Tohoku Math J* (2). 1958; 10(3):338–354.
15. Shi X, Styner M, Lieberman J, Ibrahim J, Lin W, Zhu H. Intrinsic regression models for manifold-valued data. *J Am Stat Assoc*. 2009; 5762:192–199. [PubMed: 19960103]
16. Trouné A, Vialard F-X. A second-order model for timedependent data interpolation: Splines on shape spaces. *MICCAI STIA Workshop*. 2010

Figure 1.
Schematic of the geodesic longitudinal model.

Figure 2.
Example of a Sasaki geodesic on, TS^2 , the tangent bundle of the sphere.

Figure 3. Estimation of the geodesic longitudinal model for synthetic data on the sphere. Shown are estimated individual trends (red), estimated group trend (blue), and true group trend (black).

Figure 4. Corpus callosum segmentation and boundary point model for one timepoint of one subject.

Figure 5.

Estimated corpus callosum longitudinal shape trends — (top) non-demented males and (bottom) demented males. The black shape in each figure is the estimated group intercept, $\hat{\alpha}$, representing the mean shape at baseline. The red shape in each figure is obtained by shooting along the geodesic trend determined by the estimated group slope, $\hat{\beta}$. This represents a time period of 6 years.

Table 1

Hypothesis test of group differences in corpus callosum shape changes between non-demented and demented subjects.

Variable	t^2	p -value
Intercept α	0.734	0.248
Slope β	0.887	0.027