

Chapter 24

Biomedical Applications of Geometric Functional Data Analysis



James Matuk, Shariq Mohammed, Sebastian Kurtek, and Karthik Bharath

Contents

24.1	Introduction	676
24.2	Mathematical Representation: Riemannian Metrics and Simplifying Transforms	678
24.2.1	Probability Density Functions	678
24.2.2	Amplitude and Phase in Elastic Functional Data	680
24.2.3	Shapes of Open and Closed Curves	681
24.2.4	Shapes of Surfaces.....	683
24.3	Nonparametric Metric-Based Statistics	684
24.3.1	Karcher Mean.....	684
24.3.2	Covariance Estimation and Principal Component Analysis	685
24.4	Biomedical Case Studies	686
24.4.1	Probability Density Functions	686
24.4.2	Amplitude and Phase in Elastic Functional Data	688
24.4.3	Shapes of Open and Closed Curves	692
24.4.4	Shapes of Surfaces.....	695
24.5	Summary	697
	References	698

Abstract In this chapter, we describe several biomedical applications of geometric functional data analysis methods for modeling probability density functions, amplitude and phase components in functional data, and shapes of curves and surfaces. We begin by reviewing parameterization-invariant Riemannian metrics and corresponding simplifying square-root transforms for each case. These tools

J. Matuk · S. Kurtek (✉)

Department of Statistics, The Ohio State University, Columbus, OH, USA

e-mail: matuk.3@osu.edu; kurtek.1@stat.osu.edu

S. Mohammed

Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, MI, USA

e-mail: shariqm@umich.edu

K. Bharath

School of Mathematical Sciences, University of Nottingham, Nottingham, UK

e-mail: Karthik.Bharath@nottingham.ac.uk

allow for computationally efficient implementations of statistical procedures on the appropriate representation spaces, including computation of the Karcher mean and exploration of variability via principal component analysis. We then showcase applications of these tools in multiple biomedical case studies based on various datasets including Glioblastoma Multiforme tumors, Diffusion Tensor Magnetic Resonance Image-based white matter tracts and fractional anisotropy functions, electrocardiogram signals, endometrial tissue surfaces and subcortical surfaces in the brain.

24.1 Introduction

Improvements in medical data acquisition technology, especially non-invasive imaging technology, have resulted in proliferation of large, complex datasets. There are many goals in analyzing such data depending on the application of interest, ranging from assessment of regular aging patterns to diagnosis and monitoring of various diseases. The types of imaging data of interest greatly vary in their properties, e.g., functional Magnetic Resonance Imaging (fMRI) measures dynamic brain activity through changes in blood flow, structural Magnetic Resonance Imaging (MRI) produces images of the anatomy using magnetic fields and Diffusion Tensor Magnetic Resonance Imaging (DT-MRI) maps diffusion of water molecules in biological tissues. In spite of this apparent heterogeneity, many medical imaging datasets share two common characteristics: (1) the representation space of the data is fundamentally non-Euclidean and (2) the data is functional (infinite-dimensional) in nature. These two properties of the data introduce a major challenge for statistical analysis as most traditional statistical methods apply to data residing in relatively low-dimensional Euclidean spaces. Our focus in this book chapter is on representation and statistical analysis of various aspects of biomedical imaging data including (1) patterns of voxel values via probability density functions (pdfs, smoothed histograms of voxel intensities) [44], (2) elastic functional data that contains amplitude and phase variabilities [48], (3) shapes of curves [30, 47] and (4) shapes of surfaces representing objects in medical images [19, 35]. As will be seen later, all of these data types benefit from a Riemannian geometric approach to data analysis. To unify these different data objects of interest, we refer to them commonly as geometric data objects throughout.

Statistical analysis of geometric data objects starts with the definition of a suitable mathematical representation and metric that can be used for their comparison. Once an appropriate representation space and a Riemannian metric on that space have been defined, statistical analysis proceeds via the metric structure. In particular, this approach allows one to (1) compute summary statistics such as the mean and covariance, (2) explore variability in a sample via adaptations of principal component analysis and (3) define basic statistical models [20, 49]. We consider each of pdfs, elastic functional data, and shapes of curves and surfaces separately to define the relevant Riemannian geometric representation spaces. To tie all of

the frameworks together, we point out the commonalities between the Riemannian geometry used for statistical analysis in each case.

We begin with statistical analysis of texture via a pdf representation. Texture here refers to the pattern of voxel values inside an object of interest in a medical image; it is a fundamental appearance property of objects in images [49]. We form the pdf by (1) vectorizing the relevant voxel values, (2) generating their histogram and (3) smoothing the histogram [44]. The result is a functional data object with two constraints: the pdf must be positive everywhere on its domain and it must integrate to one. The representation space of pdfs is the infinite-dimensional simplex, a constrained linear space. To define a Riemannian structure on this space, we use the well-known Fisher-Rao metric [25, 42, 46]. An important property of this metric is that it is invariant to reparameterization [7], a property used later for defining a Riemannian structure on the space of elastic functions and shapes.

The second type of geometric data objects of interest are elastic functions. Elastic functions naturally contain two different sources of variability: amplitude variability and phase, warping or parameterization variability [38]. A main goal in elastic functional data analysis is to separate these two sources of variability and define statistical methods to analyze them. The Riemannian setting for this type of analysis necessitates invariance to function reparameterization. Conveniently, we apply an extension of the Fisher-Rao metric used for pdfs in this setting [48].

Finally, we use methods from elastic shape analysis to study outlines (boundaries of objects resulting in curves and surfaces) representing objects in medical images [20, 30, 47]. The shape of such boundaries is a fundamental physical property of the objects, and provides indispensable information about the health and development of anatomical structures in the medical setting. The notion of shape is invariant to translations, scales, rotations and reparameterizations of the curves and surfaces [26]. In this setting, we use elastic Riemannian metrics which have been shown to have such desired invariances. These elastic metrics are also extensions of the Fisher-Rao metric introduced for pdfs.

In all of the above-mentioned settings, the initial Riemannian geometric structure of the representation spaces is quite complicated and necessitates numerical methods for simple tasks such as computing geodesic distances. Luckily, there exist square-root transforms in each of the cases that greatly simplify the geometry, and result in Riemannian geometric tools with analytical expressions. This, in turn, allows for development of large-scale data analytic approaches that can be applied in various biomedical settings.

Our focus in this book chapter is not on describing recent methodological advances in this area, but rather on elucidating various biomedical applications of geometric methods for functional data analysis. While we outline the relevant mathematical details to keep our discussion self-contained, the main aim is to showcase the breadth of applicability of the methods in medical imaging. As a result, our methodological descriptions are terse and avoid many technical details; we refer the interested readers to the recent books [49] and [20] for specific details. Additionally, we highlight two closely related chapters in this volume that present complementary material. In Chapter 13, the authors focus on the problem of registering different

types of functional data as well as related mathematical/statistical properties; they also present many intuitive examples to introduce this topic. In Chapter 14, the authors provide an extension of the methods described in our chapter to trajectories on general manifolds and present examples that consider multimodal data. The rest of this chapter is organized as follows. Section 24.2 describes the Riemannian geometry of representation spaces for the four geometric data objects of interest: (1) pdfs, (2) elastic functional data, (3) shapes of curves and (4) shapes of surfaces. In Sect. 24.3, we describe a general nonparametric framework, based on tools provided by the Riemannian geometric backdrops, for computing summary statistics and assessing variability in random samples. Section 24.4 discusses multiple case studies for each type of geometric data object. Here, we draw on previous studies to showcase the breadth of biomedical applications of the described methods. Finally, we close with a brief summary in Sect. 24.5.

24.2 Mathematical Representation: Riemannian Metrics and Simplifying Transforms

We begin with a brief review of the different Riemannian metrics and representations for pdfs, amplitude and phase components of elastic functional data, shapes of curves and shapes of surfaces. In each case, we highlight a particular square-root transformation, which greatly simplifies the computational implementation of the framework. For more details on these approaches, please refer to Chapters 4 (pdfs and elastic functional data), 5 and 6 (shapes of open and closed curves, respectively) in [49], and [20] (shapes of surfaces). Throughout, we use $\|\cdot\|$ and $\langle\langle\cdot, \cdot\rangle\rangle$ to denote functional norms and inner products (not necessarily \mathbb{L}^2), and $|\cdot|$ and $\langle\cdot, \cdot\rangle$ to denote the norm and inner product in a finite-dimensional Euclidean space \mathbb{R}^k .

24.2.1 Probability Density Functions

Without loss of generality, our description focuses on univariate densities on $[0, 1]$. However, the methods described here can be generalized to the multivariate setting in a straightforward manner (see Section 4 in [44] for an example). Let \mathcal{P} denote the Banach manifold of such pdfs defined as $\mathcal{P} = \{p : [0, 1] \rightarrow \mathbb{R}_+ \mid \int_0^1 p(t)dt = 1\}$. For any point $p \in \mathcal{P}$, the tangent space is defined as $T_p(\mathcal{P}) = \{v : [0, 1] \rightarrow \mathbb{R} \mid \int_0^1 v(t)dt = 0\}$; this is a vector space of all possible perturbations of the pdf p . We proceed to define a Riemannian metric on \mathcal{P} , which will be used to compute geodesic distances between two pdfs and summary statistics of samples of pdfs. The nonparametric Fisher-Rao Riemannian metric (simply FR metric hereafter), for any two tangent vectors $v_1, v_2 \in T_p(\mathcal{P})$ is defined as [25, 42, 46]

$$\langle\langle v_1, v_2 \rangle\rangle_p = \int_0^1 v_1(t)v_2(t) \frac{1}{p(t)} dt. \tag{24.1}$$

The FR metric is invariant to reparameterizations of densities [7], a nice mathematical property. One drawback of this metric is the difficulty associated with computing geodesic paths and distances due to the fact that the metric changes from point to point on the space of pdfs, requiring numerical procedures.

To simplify computation, we choose an equivalent representation of the space \mathcal{P} via the square-root density (SRD) representation [4]. Under this representation, the complicated FR metric becomes the standard \mathbb{L}^2 metric and the space of pdfs \mathcal{P} becomes the positive orthant of the unit hypersphere in \mathbb{L}^2 . In other words, we define an isometric transformation that greatly simplifies computing. The SRD is defined as a function $\psi = +\sqrt{p}$ (we omit the + sign hereafter for notational convenience). Then, the inverse mapping is unique and is simply given by $p = \psi^2$. Hence, the space of all SRDs is given by $\Psi = \{\psi : [0, 1] \rightarrow \mathbb{R}_+ \mid \int_0^1 \psi(t)^2 dt = 1\}$. The \mathbb{L}^2 Riemannian metric on Ψ is defined as $\langle\langle w_1, w_2 \rangle\rangle = \int_0^1 w_1(t)w_2(t)dt$, where $w_1, w_2 \in T_\psi(\Psi)$ and $T_\psi(\Psi) = \{w : [0, 1] \rightarrow \mathbb{R} \mid \int_0^1 \psi(t)w(t)dt = 0\}$.

As the Riemannian geometry of Ψ equipped with the \mathbb{L}^2 metric is well-known, geodesic paths and their lengths can now be computed analytically. The geodesic distance between $\psi_1, \psi_2 \in \Psi$ is simply given by

$$d(\psi_1, \psi_2) = \theta = \cos^{-1} \left(\int_0^1 \psi_1(t)\psi_2(t)dt \right). \tag{24.2}$$

The corresponding geodesic path between $\psi_1, \psi_2 \in \Psi$ is

$$\eta(\tau) = 1/\sin(\theta)\{\psi_1 \sin(\theta(1 - \tau)) + \psi_2 \sin(\tau\theta)\}, \quad \tau \in [0, 1]. \tag{24.3}$$

It is easy to see that the geodesic distance θ is bounded above by $\pi/2$. In addition to geodesic paths and distances, we often use the exponential and inverse exponential maps for computing statistical summaries of a sample of pdfs. The exponential map at a point $\psi_1 \in \Psi$, denoted by $\exp : T_{\psi_1}(\Psi) \mapsto \Psi$, is defined as

$$\exp_{\psi_1}(w) = \cos(\|w\|)\psi_1 + \sin(\|w\|)(w/\|w\|), \tag{24.4}$$

where $\|w\| = \left(\int_0^1 w(t)^2 dt \right)^{1/2}$. The inverse exponential map, denoted by $\exp_{\psi_1}^{-1} : \Psi \mapsto T_{\psi_1}(\Psi)$, is given by

$$\exp_{\psi_1}^{-1}(\psi_2) = (\theta/\sin(\theta))(\psi_2 - \psi_1 \cos(\theta)). \tag{24.5}$$

These two mappings can be used to transfer points from the nonlinear representation space Ψ to linear tangent spaces of Ψ , and vice versa.

24.2.2 Amplitude and Phase in Elastic Functional Data

One can extend the above FR metric-based framework to more general functional data. One difficulty that arises in this setting is the need for registration when comparing or modeling such observations. This is due to the fact that functional data often contains two forms of variability: amplitude and phase [38, 41, 48, 49]. Amplitude describes the vertical variability along the y -axis while phase describes the horizontal variability along the x -axis (also called domain warping), i.e., the parameterization of the functional observations. Thus, extracting phase variability from functional data through a registration procedure requires a metric that is invariant to reparameterization. As we have already established that the FR metric is invariant to reparameterizations of pdfs, we will use its extension for functional data.

We introduce some additional notation to formalize the discussion. Without loss of generality, we restrict our attention to absolutely continuous functions on the domain $[0, 1]$, and focus only on nonlinear warpings of this domain; thus, we define the function space of interest as $\mathcal{F} = \{f : [0, 1] \rightarrow \mathbb{R} \mid f \text{ is absolutely continuous}\}$. We use the set $\Gamma = \{\gamma : [0, 1] \rightarrow [0, 1] \mid \gamma(0) = 0, \gamma(1) = 1, \gamma \text{ is a diffeomorphism}\}$ to represent all possible nonlinear domain warpings. Then, for a function $f \in \mathcal{F}$, the composition $f \circ \gamma$ denotes the domain warping of f using γ , i.e., a reparameterization of the function f . To extend the FR metric for pdfs to this more general class of functions, we start with absolutely continuous functions $f : [0, 1] \rightarrow \mathbb{R}$ such that $\dot{f} > 0$; call the set of such functions \mathcal{F}_0 and let $T_f(\mathcal{F}_0)$ denote the tangent space to \mathcal{F}_0 at f . For any $f \in \mathcal{F}_0$ and $v_1, v_2 \in T_f(\mathcal{F}_0)$, the FR metric can be redefined as [48]

$$\langle\langle v_1, v_2 \rangle\rangle_f = \int_0^1 \dot{v}_1(t) \dot{v}_2(t) \frac{1}{\dot{f}(t)} dt. \quad (24.6)$$

As in the case of densities, this metric is invariant to domain warpings, $\langle\langle v_1 \circ \gamma, v_2 \circ \gamma \rangle\rangle_{f \circ \gamma} = \langle\langle v_1, v_2 \rangle\rangle_f$, for all $\gamma \in \Gamma$, $f \in \mathcal{F}_0$ and $v_1, v_2 \in T_f(\mathcal{F}_0)$, but also difficult to work with computationally.

To alleviate this issue, we define a square-root transform similar to the SRD. Define the square-root slope function (SRSF) of f as $q = \text{sign}(\dot{f})\sqrt{|\dot{f}(t)|}$. Since we have assumed $\dot{f} > 0$, the SRSF in this case simply becomes $q = \sqrt{\dot{f}}$, i.e., the square-root of an unnormalized pdf. Importantly, under the SRSF representation, the FR metric becomes the standard \mathbb{L}^2 metric. While we have so far restricted our attention to functions with positive derivative, the SRSF allows us to treat more general cases. Next, we return to the space \mathcal{F} of all absolutely continuous functions, i.e., \dot{f} is allowed to take arbitrary values including zero (when $\dot{f} = 0$, the SRSF also takes value 0). Then, using the \mathbb{L}^2 metric on the space of all SRSFs corresponding to functions in \mathcal{F} , the FR metric implicitly extends from \mathcal{F}_0 to \mathcal{F} . If the function f is absolutely continuous then the resulting SRSF is square-integrable or an element of $\mathbb{L}^2([0, 1], \mathbb{R})$ (simply \mathbb{L}^2 for brevity) [43]. The inverse mapping from an SRSF to

its corresponding function is unique up to a vertical translation. If one additionally keeps track of the starting point $f(0)$, then the mapping is unique and is given by $f(t) = f(0) + \int_0^t q(s)|q(s)|ds$. Furthermore, the SRSF of a warped function $f \circ \gamma$ is given by $(q, \gamma) = (q \circ \gamma)\sqrt{\dot{\gamma}}$.

This basic setup allows us to define amplitude and phase mathematically. The amplitude of a function remains unchanged under warping, i.e., f and $f \circ \gamma$ have the same amplitude for any $\gamma \in \Gamma$. The amplitude is thus defined as the equivalence class $[f] = \{f \circ \gamma | \gamma \in \Gamma\}$, which contains all possible domain warpings of f . The space of all amplitudes is the quotient space \mathcal{F} / Γ . In contrast to amplitude, the definition of phase is only relative. Given two functions f_1 and f_2 , the relative phase of f_2 with respect to f_1 is defined as

$$\gamma_{21} = \arg \min_{\gamma \in \Gamma} \|q_1 - (q_2 \circ \gamma)\sqrt{\dot{\gamma}}\|, \quad (24.7)$$

where q_1 and q_2 are the SRSFs of f_1 and f_2 , respectively. This minimization is usually solved using the dynamic programming algorithm [43]. The optimization problem in Eq. (24.7) is referred to as the pairwise registration of f_2 to f_1 .

Next, we focus on defining a distance for amplitude and phase components. The distance between amplitudes of two functions f_1 and f_2 is defined as

$$d_a(f_1, f_2) = d([q_1], [q_2]) = \min_{\gamma \in \Gamma} \|q_1 - (q_2 \circ \gamma)\sqrt{\dot{\gamma}}\| = \|q_1 - (q_2 \circ \gamma_{21})\sqrt{\dot{\gamma}_{21}}\|. \quad (24.8)$$

A geodesic path between two amplitude functions can then be constructed using a straight line connecting q_1 and $(q_2 \circ \gamma_{21})\sqrt{\dot{\gamma}_{21}}$. Similarly, in order to compare the phase components of the two functions f_1 and f_2 , we use the relative phase between them, γ_{21} . Then, the phase distance is defined as

$$d_p(f_1, f_2) = \cos^{-1} \left(\int_0^1 \sqrt{\dot{\gamma}_{21}(t)} dt \right). \quad (24.9)$$

This definition is based on an adaptation of the FR metric to Γ , and is measured using the SRSFs of warping functions [27]. In fact, the SRSF of any warping function is simply an SRD. Thus, the phase distance uses the SRD representation introduced earlier to compute distances between warping functions. To construct a geodesic path between two warping functions, after transforming them to their SRSFs, one can simply use Eq. (24.3).

24.2.3 Shapes of Open and Closed Curves

The extension of methods for functional data analysis to curves in higher-dimensional Euclidean spaces comes from so-called elastic shape analysis. While functional data requires invariance to reparameterization only, shape analysis

additionally requires invariance to translation, scale and rotation, also referred to as similarity shape-preserving transformations. As in the two previous sections, we begin by introducing a Riemannian metric, which is naturally invariant to all such transformations.

Let $f : \mathcal{D} \rightarrow \mathbb{R}^k$, $k > 1$ denote an absolutely continuous, parameterized curve in the Euclidean space \mathbb{R}^k with the domain of parameterization given by $\mathcal{D} = [0, 1]$ for open curves and $\mathcal{D} = \mathbb{S}^1$ for closed curves. With a slight abuse of notation, let \mathcal{F} denote the set of all such curves. While the framework described here applies to k -dimensional curves, biomedical applications generally consider 2D and 3D curves as data objects, as seen in later sections. The most difficult of the aforementioned invariances is that to parameterization, and it requires the definition of a nonstandard Riemannian metric on \mathcal{F} referred to as the elastic metric. We begin by identifying the curve f with the pair (r, θ) where $r = |\dot{f}|$ is the speed function and $\theta = \frac{f}{|\dot{f}|}$ is the angle function. The only information lost when passing from f to the pair (r, θ) is translation, which is one of the nuisance, shape-preserving transformations. Also, let $(\delta r_1, \delta \theta_1)$ and $(\delta r_2, \delta \theta_2)$ be two tangent vectors at (r, θ) . Then, the elastic Riemannian metric is defined as

$$\langle\langle (\delta r_1, \delta \theta_1), (\delta r_2, \delta \theta_2) \rangle\rangle_{(r, \theta)} = a \int_{\mathcal{D}} \delta r_1(t) \delta r_2(t) \frac{1}{r(t)} dt + b \int_{\mathcal{D}} \delta \theta_1(t)^T \delta \theta_2(t) r(t) dt. \tag{24.10}$$

We note three important properties of this metric. First, it is a weighted combination of two terms, one capturing changes in the speed function, i.e., stretching deformations, and one capturing changes in the angle function, i.e., bending deformations. Second, the stretching term in the metric should look familiar: it is the same as the FR metric introduced earlier for densities. Third, this metric is invariant to reparameterizations of curves, in addition to translation, scaling and rotation. Unfortunately, as in the two previous cases, this metric is difficult to use in practice.

Fortunately, one can extend the SRSF representation introduced for functional data to this more general case. This new representation of curves is called the square-root velocity function (SRVF) [21] and is defined as $q = \sqrt{r} \theta = \frac{f}{\sqrt{|\dot{f}|}}$. In fact, the SRVF and SRSF are equivalent for univariate curves. The SRVFs of absolutely continuous curves reside in $\mathbb{L}^2(\mathcal{D}, \mathbb{R}^k)$ (simply \mathbb{L}^2 for brevity). An important property of this representation is that the complicated elastic metric, with $a = 1/4$ and $b = 1$, simplifies to the standard \mathbb{L}^2 metric under the SRVF transform. We note that the SRVF is not the only transform that simplifies a specific instance of the elastic metric to the \mathbb{L}^2 metric; for alternative approaches see [2, 28, 53, 54]. We will use the SRVF to mathematically formalize the notion of shape so that any two curves that are within a translation, rescaling, rotation and reparameterization of each other are considered to be the same data object. Since the SRVF is a function of the derivative of the original curve, it is automatically translation invariant (this is obvious since the elastic metric is translation invariant). Forcing a unit length constraint on the curves results in unit \mathbb{L}^2 norm SRVFs, i.e., $\|q\|^2 = 1$. Hence, the set of unit length open curves is $C = \{q : [0, 1] \rightarrow \mathbb{R}^k \mid \|q\|^2 = 1\}$, i.e., a unit sphere in \mathbb{L}^2 ; C is

also referred to as the pre-shape space. Restricting attention to closed curves, the pre-shape space becomes $C^c = \{q : \mathbb{S}^1 \rightarrow \mathbb{R}^k \mid \|q\|^2 = 1, \int_{\mathbb{S}^1} q(t)|q(t)|dt = 0\}$, which is a subspace of C due to the closure constraint. In the remainder, to keep the discussion general, we do not make a distinction between these two pre-shape spaces and simply use C . The rotation and reparameterization variabilities can be filtered out through a suitable definition of equivalence classes. Let $[q] = \{O(q, \gamma) \mid \gamma \in \Gamma, O \in SO(k)\}$ denote all possible rotations and reparameterizations of q , where $SO(k) = \{O \in \mathbb{R}^{k \times k} \mid O^T O = O O^T = 1, \det(O) = 1\}$ is the special orthogonal group of rotations, $\Gamma = \{\gamma : \mathcal{D} \rightarrow \mathcal{D} \mid \gamma \text{ is a diffeomorphism}\}$ is the set of (order-preserving) reparameterizations and $(q, \gamma) = (q \circ \gamma)\sqrt{\dot{\gamma}}$. Each equivalence class represents a shape uniquely and the collection of all equivalence classes is the shape space $\mathcal{S} = C/(SO(k) \times \Gamma)$. The final ingredient is the ability to compare shapes using a distance on \mathcal{S} . Under the SRVF representation, this distance is given by

$$d([q_1], [q_2]) = \min_{O \in SO(k), \gamma \in \Gamma} \cos^{-1} \left(\int_{\mathcal{D}} q_1(t)^T O q_2(\gamma(t)) \sqrt{\dot{\gamma}(t)} dt \right). \quad (24.11)$$

The optimization problem in Eq. (24.11) is solved using a combination of Procrustes analysis [10] and dynamic programming [43]. For visualization, a geodesic path between two shapes can be constructed using Eq. (24.3) with inputs q_1 and $O^*(q_2, \gamma^*)$, where O^* and γ^* denote the minimizers of Eq. (24.11).

24.2.4 Shapes of Surfaces

Lastly, we consider shape analysis of surfaces. This case evolves similarly to the case of curves. Again, with a slight abuse in notation, let \mathcal{F} denote the space of smooth embeddings $f : \mathcal{D} \rightarrow \mathbb{R}^3$, where the domain of parameterization \mathcal{D} can be a unit sphere (closed surfaces), a unit square (quadrilateral surfaces), a unit cylinder (cylindrical surfaces), a unit disk (hemispherical surfaces), etc. Furthermore, let Γ be the set of all diffeomorphisms of \mathcal{D} . We use $n(t) \in \mathbb{R}^3$ to denote the normal vector to the surface at the point $t \in \mathcal{D}$, i.e., $n(t) = \frac{\partial f}{\partial u}(t) \times \frac{\partial f}{\partial v}(t)$, where (u, v) are the coordinates on the domain \mathcal{D} . The infinitesimal area measure at a point t is given by $r(t) = |n(t)|$ and the normalized normal vector is $\tilde{n}(t) = \frac{n(t)}{r(t)}$. We will represent the surface f using the pair (r, \tilde{n}) ; as this representation depends on partial derivatives only, it is automatically invariant to translations. Let $(\delta r_1, \delta \tilde{n}_1)$ and $(\delta r_2, \delta \tilde{n}_2)$ be two tangent vectors at (r, \tilde{n}) . A reparameterization invariant Riemannian metric on the space of surfaces is given by [19]

$$\langle\langle (\delta r_1, \delta \tilde{n}_1), (\delta r_2, \delta \tilde{n}_2) \rangle\rangle_{(r, \tilde{n})} = \frac{1}{4} \int_{\mathcal{D}} \delta r_1(t) \delta r_2(t) \frac{1}{r(t)} dt + \int_{\mathcal{D}} \delta \tilde{n}_1(t)^T \delta \tilde{n}_2(t) r(t) dt. \quad (24.12)$$

Again, the first term in this metric resembles the FR metric introduced earlier, and captures changes in the infinitesimal areas of surface patches, i.e., stretching deformations. The second term captures changes in the direction of the unit normal

vector, i.e., bending deformations. The metric in Eq. (24.12) is a special case of a more general elastic metric for surfaces [19]. Due to the difficulty of working with this metric in practice, we define an alternative representation of surfaces, called the square-root normal field (SRNF), which simplifies this metric to the standard \mathbb{L}^2 metric. The SRNF of a surface f is given by $q = \sqrt{r}\tilde{n} = \frac{n}{\sqrt{|n|}}$. The SRNF of a reparameterized surface $f \circ \gamma$, for a $\gamma \in \Gamma$, is given by $(q, \gamma) = (q \circ \gamma)\sqrt{J_\gamma}$, where J_γ is the determinant of the Jacobian of γ .

As in the case of curves, we seek a framework that is invariant to all shape-preserving transformations (translation, scale, rotation and reparameterization). The SRNF representation is automatically invariant to translations. To produce invariance to scaling, we rescale all surfaces to unit area, resulting in SRNFs with unit \mathbb{L}^2 norm. As in the case of curves, this amounts to restricting attention to the unit sphere in \mathbb{L}^2 . We then define a distance on the shape space of surfaces by minimizing over equivalence classes of the form $[q] = \{O(q, \gamma) | \gamma \in \Gamma, O \in SO(3)\}$

$$d([q_1], [q_2]) = \min_{O \in SO(3), \gamma \in \Gamma} \cos^{-1} \left(\int_{\mathcal{D}} q_1(t)^T O q_2(\gamma(t)) \sqrt{J_\gamma(t)} dt \right). \quad (24.13)$$

As in the case of curves, the optimal rotation is found using Procrustes analysis [10]. Computation of the optimal reparameterization requires a gradient descent algorithm [29]. A geodesic path between two shapes can be constructed using Eq. (24.3) with inputs q_1 and $O^*(q_2, \gamma^*)$, where O^* and γ^* are the minimizers of Eq. (24.13).

24.3 Nonparametric Metric-Based Statistics

We provide a general recipe for computing the sample mean, covariance and performing principal component analysis (PCA). Our tools rely on Karcher means for metric spaces and local linear approximations via the Riemannian structure. Since all four geometric data objects described in Sect. 24.2 rely on \mathbb{L}^2 Riemannian geometry, we provide a single description here for brevity.

24.3.1 Karcher Mean

The sample Karcher mean [24] of a collection of points (i.e., pdfs, amplitude functions, phase functions or shapes) x_1, \dots, x_n from a metric space (\mathfrak{X}, d) is defined as the minimizer of the Karcher variance

$$\hat{\mu} = \arg \min_{x \in \mathfrak{X}} \frac{1}{n} \sum_{i=1}^n d(x, x_i)^2. \quad (24.14)$$

This definition, with slight modification when dealing with equivalence classes, is applicable to all four metric spaces discussed in Sect. 24.2. Computation of the Karcher mean is carried out using gradient-based algorithms [31, 36, 40], which generally iterate between three steps: (1) projection of data from the representation space to the linear tangent space at the current estimate of the mean via the inverse exponential map, (2) computation of the gradient of the cost function in Eq. (24.14), and (3) update of the current estimate of the mean using the exponential map. In the case of functional data, the Karcher mean is used as a template for multiple function registration. That is, once the Karcher mean is estimated, the amplitude components of all functions are defined through pairwise registration to the Karcher mean; this also results in the phase component, computed with respect to the mean [48].

24.3.2 Covariance Estimation and Principal Component Analysis

Exploration of variability in a sample of geometric data objects can be carried out by choosing local coordinates in the vicinity of the Karcher mean $\hat{\mu}$. The Riemannian structure allows one to conveniently linearize the data representation space via the tangent space at the mean, $T_{\hat{\mu}}$, and to select Euclidean coordinates in this space.

As before, let x_1, \dots, x_n and $\hat{\mu}$ represent the data objects of interest and their Karcher mean, respectively. We begin by projecting each x_i , $i = 1, \dots, n$ onto the tangent space at the mean using the inverse exponential map resulting in tangent vectors v_1, \dots, v_n . Using this tangent space representation, we estimate the covariance matrix based on discretized versions of the tangent vectors denoted by \mathbf{v}_i , $i = 1, \dots, n$. Assuming the dimension of each \mathbf{v}_i is M , the sample covariance matrix is given by $K_M := 1/(n-1) \sum_{i=1}^n \mathbf{v}_i \mathbf{v}_i^T$. To study variability using PCA, we apply the spectral decomposition to the covariance matrix $K_M = U \Sigma U^T$, where the orthogonal matrix U contains the principal components (PCs) or principal directions of variability, and the diagonal matrix Σ contains the PC variances. In typical biomedical applications, the number of observations is smaller than the dimensionality of each tangent vector, i.e., $n < M$. Thus, there are at most $n - 1$ positive values in the matrix Σ . The submatrix formed by the first r columns of U , U_r , spans the r -dimensional principal subspace of the observed data, and one can reexpress the data using coordinates in this subspace via principal coefficients computed as $c_i = U_r^T \mathbf{v}_i$, $i = 1, \dots, n$. One can then use these principal coefficients for further statistical modeling, e.g., PC regression [3]. A common approach to modeling complex data objects is through tangent PCA-based models such as the truncated wrapped Gaussian distribution [30] or by directly modeling the principal coefficients.

24.4 Biomedical Case Studies

We focus on multiple biomedical case studies that consider (1) pdfs, (2) amplitude and phase in functional data, (3) shapes of curves, and (4) shapes of surfaces as data objects. While the theoretical underpinnings outlined in Sect. 24.2 consider infinite-dimensional data representations, computer implementation of these methods requires appropriate discretization. We represent pdfs and other univariate functions (amplitude/phase) as $1 \times N$ vectors, where N denotes the number of discretization points selected on the function domain. Shapes of curves are represented as $d \times N$ matrices, where $d = 2, 3$ depending on whether the curve is planar or 3D, and N is again the number of points selected on the curve domain. Finally, shapes of surfaces are represented as $N_1 \times N_2 \times 3$ arrays, where $N_1 \times N_2$ defines a discretization grid on the surface, and each point on the grid takes a value in \mathbb{R}^3 .

24.4.1 Probability Density Functions

Assessment of Glioblastoma Multiforme Tumor Texture Variability Glioblastoma multiforme (GBM), also known as grade IV glioma, is the most common form of a malignant brain tumor in adults [15]. It is a morphologically heterogeneous disease with extremely poor prognosis; also, predicting the impact of standard cancer treatments such as chemotherapy and radiation therapy becomes considerably challenging. Thus, exploring tumor heterogeneity is critical in cancer research as inter- and intra-tumor differences have stymied the systematic development of targeted cancer therapies [9]. MRI is one of the modern medical imaging techniques that has been used to investigate tumor development in various contexts. MRI scans are primarily used to exhibit and evaluate the location, size, growth and progression of tumors, which serve as indicators for clinical decision making. Various physiological features are extracted by using voxel-level data to visualize the progression (or regression) of tumors. This is generally done by constructing voxel value histograms. However, in most cases, only simple summaries of the entire histograms are used for statistical analysis. This approach has two main drawbacks. First is the subjectivity in the choice of the number and location of the summary features (e.g., quantiles or percentiles, etc.). Second, and more importantly, these summary features fail to capture the entire information in a histogram of voxel intensities, and thus cannot detect small-scale and sensitive changes in the tumor due to treatment effects [23].

Alternatively, one can exploit the entire histogram, or its corresponding smoothed density profile, for the tumor region in an MRI. This was the approach taken in a recent paper that introduces DEMARCATE, a self-contained pipeline for geometric clustering and validation of GBM tumor texture profiles [44]. Semi-automated segmentation methods [1] can be employed to delineate the tumor region in the

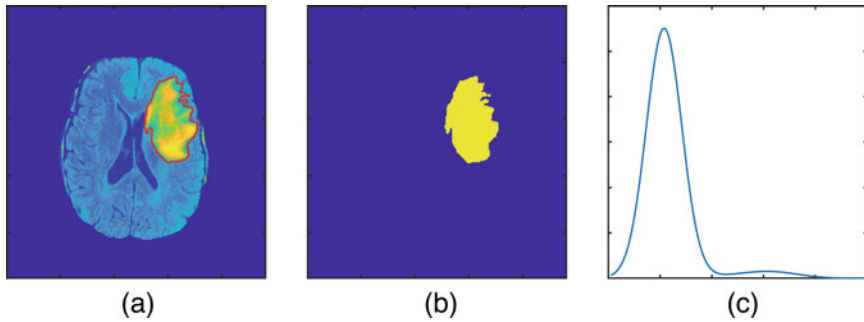


Fig. 24.1 (a) MRI slice for a subject with GBM; the delineated region corresponds to the tumor. (b) Mask identifying the tumor region. (c) Estimated voxel intensity pdf corresponding to the tumor

whole brain MRI scan. In subsequent analyses we use the voxel-level information from the axial slice with the largest tumor area only. This is done for simplicity of visualization and can be easily extended to the full 3D tumor. Figure 24.1a shows a single slice of an MRI scan for a subject with GBM, where the delineated region corresponds to the tumor. This region is displayed as a binary mask in panel (b). The voxel values inside the tumor are used to compute a Gaussian kernel density estimate (a pdf), which is displayed in panel (c); it contains detailed and refined information about the voxel-level tumor characteristics. Hence, under this setup, a sample of GBM scans is represented by a sample of voxel value pdfs corresponding to the tumor region in the MRI scan of each subject. For a more detailed description of the image processing pipeline, we refer the interested reader to [44]. The imaging data used in this study was retrieved from The Cancer Imaging Archive (www.cancerimagingarchive.net).

Next, we consider a comparison of two subjects based on their voxel value pdfs. Figure 24.2a, b shows the MRI slice for two subjects, and the corresponding pdfs of the tumor intensity values. The geodesic path between the two pdfs under the FR metric is shown in Fig. 24.2c. The displayed geodesic was discretized with five equally spaced points on the interior of the path. Finally, we consider a random sample of ten subjects with GBM. The densities for these ten subjects (dashed), along with their Karcher mean (solid red) are displayed in Fig. 24.3a. The Karcher mean in this case provides a simple summary of the sample of voxel intensity pdfs, and was computed using the FR Riemannian framework. We do not display the corresponding MRI slices in this case for brevity (note that there doesn't exist a unique MRI slice corresponding to the Karcher mean pdf). Given an estimate of the Karcher mean, we perform PCA and show the first principal direction of variability in the given sample. This result is provided in Fig. 24.3b and reflects the relative heights of the different modes in the sample of voxel value pdfs. While not shown here, principal coefficients can be subsequently used as covariates in regression models, e.g., to predict subject survival [3].

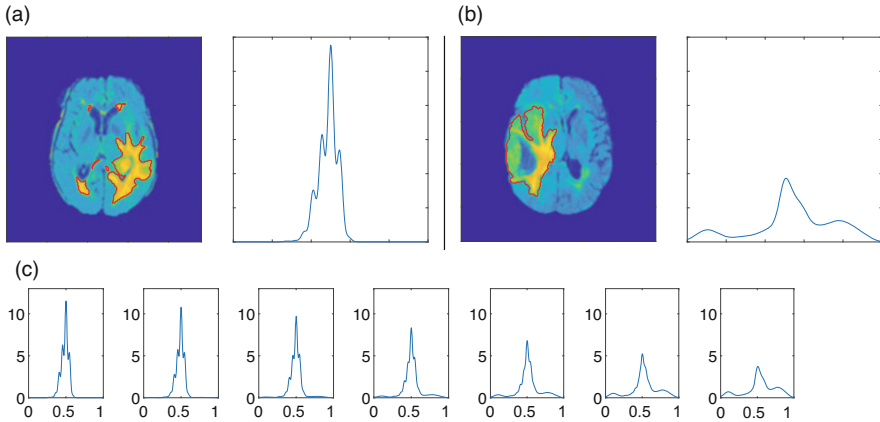


Fig. 24.2 (a and b) MRI slices from two different GBM subjects, with the pdf corresponding to the tumor intensity values. (c) Geodesic path between the pdfs for subject (a) and subject (b)




Fig. 24.3 (a) Karcher mean (solid red line) of a random sample of ten voxel value pdfs (dashed lines) extracted from tumor regions of GBM subjects. (b) Principal direction of variability in the sample, displayed at -2 , -1 , 0 , $+1$, $+2$ standard deviations around the mean (red)

24.4.2 Amplitude and Phase in Elastic Functional Data

Automatic Segmentation and Clustering of Electrocardiogram Signals The electrocardiogram (ECG) is a cheap and widely-applied diagnostic tool for assessment of various heart diseases including myocardial infarction (MI). Automated algorithms, based on sound mathematical and statistical principles, that can accurately and efficiently analyze ECG signals are thus useful in monitoring and identifying the risk or onset of a particular disease. The ECG captures fluctuations in electrical potential of the heart muscle on the body surface and results in a vector that represents the magnitude and direction of the electric field generated through the heart [8]. The ECG represents an example of a highly periodic biomedical signal. The two main challenges in analyzing such data include (1) automatic segmentation

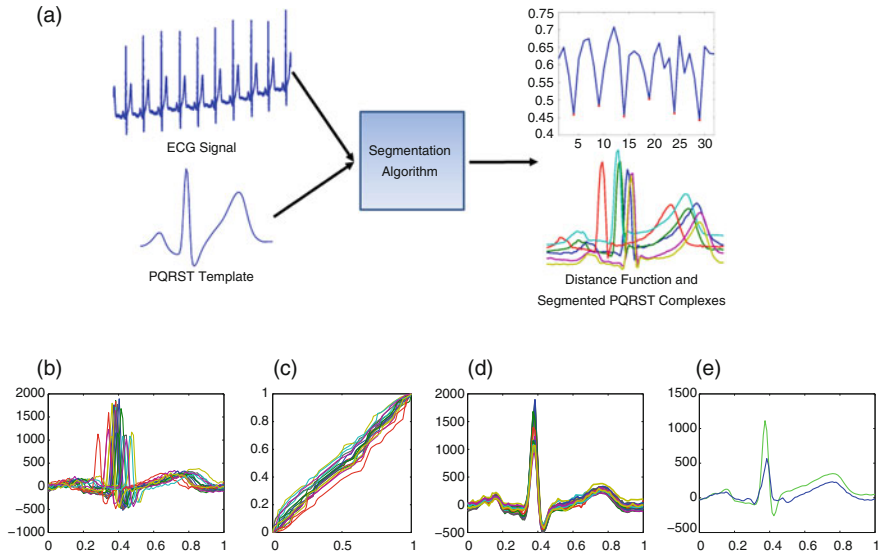


Fig. 24.4 (a) Pictorial description of automatic algorithm for segmentation of long ECG signals. Bottom: Registration of PQRST complexes to a common template. (b) Given PQRST complexes. (c) Phase component. (d) Amplitude component. (e) Comparison of average PQRST complex without (blue) and with (green) registration. Image courtesy of [32]

of cycles called PQRST complexes (PQRST refer to semantic features of each cycle: the P peak, QRS complex and the T peak) from a long temporal ECG signal, and (2) automatic registration of cycles to extract amplitude and phase variabilities of individual cycles. The ECG data used here for demonstrative purposes is a subset of the PTB Diagnostic ECG Database [5] obtained from Physionet [12].

In [32], the authors solve these two problems using techniques from elastic functional data analysis described in Sect. 24.2.2. First, they define an automatic signal segmentation algorithm based on a sliding window approach. In particular, they construct a PQRST complex template, based on the amplitude component of a few manually segmented PQRST complexes, and slide it along the long periodic signal. The cost function that is then used for segmentation is the phase distance, defined in Eq. (24.9), between the part of the long signal in the current window and the defined template. The PQRST cycles are identified as local minima of this cost function. Figure 24.4a provides a pictorial description of this process. Once the cycles have been extracted, their amplitude components are found by registering to a new common template. This result is displayed in the bottom panel of Fig. 24.4. In (b), we show the segmented PQRST complexes. The extracted phase components are displayed in (c) with corresponding amplitude components in (d). Finally, in (e), we compare the amplitude means computed without (blue) and with (green) registration. Note the enhanced features of the PQRST complex average computed after registration.

In addition to extracting the amplitude and phase components from PQRST complexes, the authors in [32] use the amplitude components for (1) classification of subjects as healthy controls or as having MI and (2) localization of the MI as anterior or inferior. The data they use for this experiment consists of 80 healthy control ECGs, 28 of which are repeated measures for the same subject, and 80 MI ECGs with no repeated measures. For each subject, they first segment the long ECG signal into corresponding PQRST complexes and then use the amplitude of the average PQRST complex for classification using the nearest neighbor procedure. They report an accuracy of MI classification of 90% by combining information from different ECG leads (the data contains a total of 15 different ECG signals called leads per subject). Also, they report a localization accuracy of 92.21%, again based on combining multiple single lead classifiers.

Assessment of Variability in DT-MRI Fractional Anisotropy Functions in Multiple Sclerosis DT-MRI is a neuroimaging modality that traces the diffusion of water molecules in the brain. A DT-MRI scan of a subject's brain provides a 3×3 tensor matrix, at each voxel in the image, that describes the constraints of local Brownian motion of water molecules. This information is essential to understanding white matter in the brain which constitutes areas made up of axons or tracts. Tracts connect neurons and allow for the transmittance of electric signals from one area of the brain to another, affecting overall brain function. Because the diffusion of water in tracts is anisotropic, tracts themselves can be extracted from the information contained in a DT-MRI, along with other quantities of interest that describe the quality of tract connection by summarizing its degree of anisotropy.

Here, we focus on Fractional Anisotropy (FA) measurements along tracts, which provide a voxel-wise summary of the eigenvalues of the diffusion tensors, denoted by $\lambda_1, \lambda_2, \lambda_3$. At each voxel, FA is given by the scalar quantity $FA = \sqrt{\frac{3}{2} \frac{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}$, where $\bar{\lambda} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}$. A larger FA value indicates a large degree of anisotropy. For practitioners, this summary is interpreted as measuring the quality of connections between neurons connected by the tracts in a particular region of interest, and has been found to be a useful quantity to study subjects with various diseases, e.g., multiple sclerosis (MS) [13], Alzheimer's [39], etc. In the MS setting, the autoimmune disease causes lesions and damage to tracts that results in a decrease in FA. Thus, this quantity can be used to distinguish between healthy controls and subjects with MS, and to predict cognitive and motor disease outcomes. The data of interest takes a functional form, with the domain of the functions representing locations along tracts. Determining the voxels that the tracts pass through in the image is a practical challenge in itself and will be further discussed in Sect. 24.4.3.

The functional FA data we analyze here is available as part of the 'refund' package in R [14]. In particular, we study the mean and principal directions of variability in a sample of 66 subjects with MS whose FA values were measured at 55 locations along the right corticospinal tract that contributes to fine motor movements in ipsilateral limbs. The domain of parameterization for each FA function was




Fig. 24.5 (a) Observed FA functions. (b) Amplitude component. (c) Phase component

normalized to $[0, 1]$ for convenience. It is important to note that due to differences in the geometry of different subjects' white matter, there generally exist both phase and amplitude variabilities in the FA functional data, as demonstrated next. The raw FA functions for the 66 subjects are shown in Fig. 24.5a. The amplitude components of the functions, after registration to a common template, are displayed in Fig. 24.5b. Finally, the warping functions which constitute the phase components are shown in Fig. 24.5c. Visual inspection of panel (b) reveals that the number of extreme values in the FA functions is roughly the same across subjects. The main source of variability in this case are the heights of the extreme values. The phase components in panel (c) suggest that the extreme values occur at different parameter values for different subjects, which is intuitive given natural geometric variability of the tracts across subjects. These insights are only made possible through the separation of amplitude and phase by registering all functions to a common template; such patterns are much more difficult to observe by looking at the observed functions in panel (a).

Figure 24.6 further highlights the importance of elastic functional data analysis methods by contrasting averages computed without (panel (a)) and with (panel (b)) registration. While the general patterns in the two means are similar, the amplitude mean in panel (b) reveals much more local structure through small peaks and valleys. Finally, to understand patterns of variability in the given sample of FA functions, we perform PCA on the amplitude components. Since the translation of the functions is also informative in this setting, we include it as an additional feature in the PCA model (it is appropriately weighted to make the scales of the two components, amplitude and translation, comparable). The first three principal directions of amplitude (and translation) variability are visualized in Fig. 24.6c–e. The first direction predominantly captures variability in translation as well as the initial portion of the functions, as some functions in the sample initially decrease and others increase. The second direction captures fine features of the different peaks and valleys of the FA functions, especially the fourth peak, as well as large amount of variability at the end of the functions. Finally, the third direction (and subsequent directions not displayed here) capture bigger differences in the relative heights of peaks and valleys.

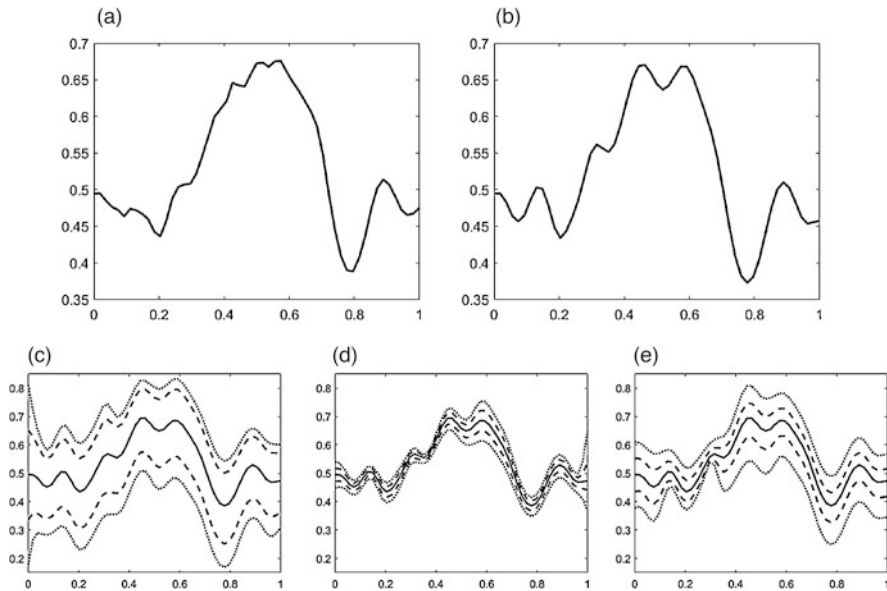


Fig. 24.6 (a) Pointwise mean of the FA functions in Fig. 24.5a. (b) Karcher mean of the FA functions after registration in Fig. 24.5b. (c)–(e) First three principal directions of amplitude (and translation) variability of the FA functions, respectively. We display a path of functions sampled at -2 and $+2$ (dotted lines), -1 and $+1$ (dashed lines), and 0 (solid line) standard deviations from the mean

24.4.3 Shapes of Open and Closed Curves

Comparison and Summarization of Planar Shapes of GBM Tumors We return to the study of the GBM tumor dataset, as described in Sect. 24.4.1. However, instead of modeling the internal texture of the tumors, we instead model the shapes of tumor outlines. This allows us to study growth patterns and shape heterogeneity of tumors, which are features that are complementary to voxel intensity values. Tumor shape is affected by the location of the tumor in the brain due to constraints posed by the brain anatomy such as white matter and blood vessels. In [3], the authors suggest that tumor shape could enhance our understanding of disease prognosis and help in prediction of therapeutic success. As in Sect. 24.4.1, the imaging data is a subset of The Cancer Imaging Archive, and the tumor shape is obtained through semi-automated segmentation; a segmented tumor is visualized in Fig. 24.1. The geometric data object of interest in this case is the red outline of the tumor rather than the entire MRI slice. In this case study, we consider 63 GBM tumor outlines, which are represented as planar closed curves. We focus on characterization of tumors through the visualization of geodesic paths, the Karcher mean shape and shape PCA. Similar results appear in [3]; the scope of their study is broader and additionally includes shape clustering, hypothesis testing and survival modeling.

We begin with a visualization of a geodesic path between two tumor shapes in Fig. 24.7. If the two endpoints of the geodesic path are a single subject’s tumor at different timepoints, the points along the path can be viewed as an interpolation along different stages of tumor growth. This, in turn, can help a practitioner retrospectively understand how the subject’s tumor has evolved over time without collecting MRI data at intermediate timepoints. On the other hand, when the two endpoints are shapes of tumors coming from two different subjects, as in Fig. 24.7, the path can help formulate a qualitative understanding of how tumor shapes differ in the population. In this case, the shapes of the tumors seem to differ by how bulbous or skinny their protrusions are. By viewing more subjects’ tumors in Fig. 24.8, this seems to be a common discrepancy between the different subjects. The insight that this is a primary source of variability in GBM tumor shapes is formalized by viewing the principal directions of variability in the entire dataset; the first four directions are shown in Fig. 24.9. Notice that the first direction, which captures approximately 41% of the total variability, describes the types of differences in




Fig. 24.7 Geodesic path between shapes of GBM tumors for two subjects (blue and black endpoints), sampled uniformly using five interior points along the path

Fig. 24.8 Five randomly selected GBM tumor outlines



Fig. 24.9 (a)–(d) First four principal directions of shape variability in the GBM dataset, respectively, sampled at $-3, -2, -1, 0, +1, +2$ and $+3$ standard deviations around the mean shape (red)



protrusions described before. The remaining directions describe other shapes of the bulbous features of the tumors. The second, third and fourth principal directions of variation capture about 33%, 16%, and 10% of the total variability, respectively; essentially all of the variation is contained in these first four directions. This implies that a low dimensional model, based on these PCs, could be used for subsequent statistical analyses.

Clustering Shapes of 3D DT-MRI Tracts As previously mentioned in Sect. 24.4.2, DT-MRI tracts are of interest when studying structural connections between different regions of the brain. Tractography is the field of study concerned with discerning tracts from the tensor-based DT-MRIs [37]. Conventional tractography relies on the principle that water diffuses anisotropically in white matter tracts in a principal direction that is encoded in the diffusion tensor. This implies that the direction that the tract is pointing in a voxel will coincide with the eigenvector corresponding to the maximum eigenvalue of the diffusion tensor. Consequently, an entire tract can be traced using information from the observed diffusion tensors associated with voxel locations. The application described in [30] deals with tracts that connect Broca's and Wernicke's regions of the left hemisphere of the brain; these two regions are associated with the human language circuit. While two main routes of connection are widely recognized, there is an ongoing debate on whether the white matter tracts connecting the two regions can be further broken down into smaller subroutes.

The data in this study contains different numbers of fiber tracts for four subjects. We identify different routes of connectivity by clustering the shapes of these tracts using distance-based methods. This was also done in [30], but there the authors used shape in conjunction with other features of the tracts. To determine if the tracts can be put in different clusters representing major pathways connecting the regions of interest, a hierarchical clustering algorithm, with a complete linkage criterion, is used to cluster the observations for each individual based on the elastic shape distance defined in Eq. (24.11). The results for all four subjects are shown in Fig. 24.10. The tracts, represented as 3D open curves, are plotted in the top panel of the figure and are colored by cluster membership. In the middle panel, we show the pairwise shape distance matrix as an image, rearranged according to the computed clusters. Note the nice separation of clusters in this distance matrix. Finally, in the bottom panel, we show a plot of the tracts after applying multidimensional scaling (MDS) to the distance matrices. This 2D scatterplot provides a lower dimensional visualization of the clustered data. Some of the subjects exhibited tracts that could be separated into more than two clusters, e.g., Fig. 24.10b, c. The case for more than two clusters is hard to justify for the subject in Fig. 24.10a. Based on these results, it appears that the hypothesis that there are two or three main pathways connecting Broca's and Wernicke's regions in the left hemisphere is plausible for all of the subjects considered in this case study.

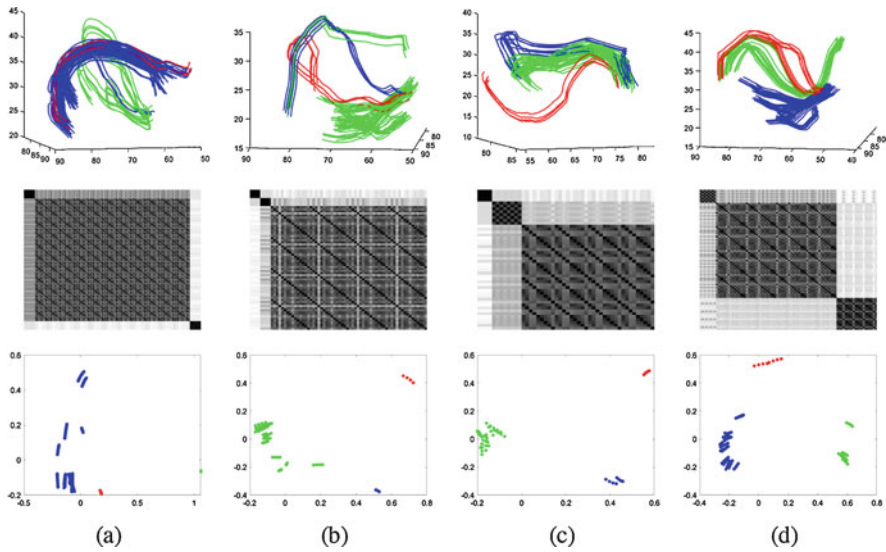



Fig. 24.10 (a)–(d) Results of hierarchical shape clustering for four subjects. Top: Tracts colored by cluster membership. Middle: Image of distance matrix. Bottom: MDS plot of tracts

24.4.4 Shapes of Surfaces

Simulation of Endometrial Tissue Shapes We define a PCA-based statistical model for efficient simulation of random endometrial tissue shapes, which can be used for validation of various image processing algorithms such as multimodal registration of MRI and transvaginal ultrasound (TVUS). This is an important task in the context of diagnosis and surgery planning for endometriosis [45, 52], a complex gynecological disease in which endometrial cells appear outside their usual locations in the uterine cavity [6]. Endometriosis affects approximately 10% of women in the reproductive age group and may cause chronic pelvic pain, severe dysmenorrhea, infertility, rectal bleeding and digestive problems.

In this study [33, 34], we use real data from ten subjects with small endometrial implants in the pelvic area. The available data are cylindrically parameterized surfaces of endometrial tissues, reconstructed from 2D MRI slices. The entire dataset can be found in Figure 1(b) in [34]. There is a lot of variation in this data, and thus, parsimonious shape models are very important in this application. Of main interest is random generation of realistic endometrial tissue shapes as they'd appear in an MRI scan and a corresponding TVUS image. Unfortunately, endometrial tissue is soft and undergoes a significant deformation during TVUS imaging due to the transducer's pressure. Thus, in addition to generating a random endometrial tissue shape we must additionally apply a deformation on the surface of the shape that is consistent with the TVUS imaging protocol.

Fig. 24.11 (a) Randomly sampled shape from the Gaussian model resembling MRI data. (b) Random sample after additional deformation resembling TVUS data. (c) Deformation applied to the random sample displayed on a perfect cylinder. Image courtesy of [34]



To achieve the two goals outlined above, we first compute the Karcher mean of and perform PCA on the ten given endometrial tissue shapes. This allows us to express the data in terms of the principal coefficients. We model these coefficients using a simple zero-mean multivariate Gaussian distribution with the covariance structure informed by the PCA. A major advantage of this shape model is that it is very easy and computationally efficient to sample from. Figure 24.11a shows four randomly generated endometrial tissue shapes as they'd appear in an MRI. Then, to simulate the semi-synthetic deformation needed for the corresponding TVUS-based endometrial tissue shape, we define a simple diffusion model with different degrees of deformation on the previously computed Karcher mean; the deformation is centered at a randomly selected point on the mean. These deformations can then be transported from the Karcher mean to each of the random samples from our model using parallel transport [51]. Figure 24.11c displays the deformations applied to a perfect cylinder. The magnitude of deformation increases from the top row to the bottom row. Finally, the TVUS-based, deformed endometrial tissue shapes are displayed in (b). The random samples generated using this approach (as well as their deformed counterparts) naturally resemble the given data. In [34], the authors provide a thorough validation of their models and a formal assessment by a clinician.

Classification of Attention Deficit Hyperactivity Disorder (ADHD) via Shapes of Subcortical Structures Recently, many researchers have become interested in studying shape changes of brain structures and associating these changes with various diseases including Alzheimer's [22, 50], Parkinson's [11], autism [16] and ADHD [29], among others. Statistical analysis of the shapes of such structures plays a central role in the ability to diagnose and monitor such diseases, as well as to develop novel treatment strategies. The current standard of practice is to use clinical




Fig. 24.12 Subcortical structures used for classification of ADHD. Image courtesy of [20]

symptoms, including various behavioral tests, to detect and quantify abnormalities due to disease status. Such an approach has clear limitations as the tests are often subjective and mainly qualitative, relying entirely on a doctor's assessment and judgment.

As an alternative, our final case study considers classification of ADHD based on the shapes of four distinct subcortical structures, represented as closed surfaces: pallidum, caudate, thalamus and putamen; a single example of each structure is displayed in Fig. 24.12. The surfaces of these subcortical structures were segmented from T1-weighted MRIs of young adults aged between 18 and 21 who were recruited from the Detroit Fetal Alcohol and Drug Exposure Cohort [17, 18]. Among the 34 subjects in this dataset, 19 were diagnosed with ADHD and the remaining 15 were controls. The classifier in this study was constructed in the following way. First, the training data was used to estimate the Karcher mean in each class. Then, shape PCA was used to define a Gaussian model on the principal coefficients. The resulting classifier simply uses the likelihood ratio under these two models to classify test shapes into control or ADHD classes. This classifier was applied in a leave-one-out manner to the above-described dataset, i.e., at each iteration a single case was left out for testing while the rest were used to learn the classification model. The best classification result obtained using this method was based on the shape of the left putamen, 94.1%. The shape of the right pallidum yielded a classification accuracy of 76.5%, and the shapes of the left caudate, left thalamus and right thalamus resulted in a slightly worse classification accuracy of 67.7%. Comprehensive results of this study are reported in [20], where the approach outlined here was compared to other classifiers and other shape representations [29].

24.5 Summary

We consider several biomedical applications of geometric functional data analysis. We begin by assessing variability in a sample of GBM voxel intensity pdfs to model tumor appearance. We then shift our focus to the use of elastic functional data methods for analyzing amplitude and phase components of electrocardiogram signals and FA functions extracted from DT-MRI. For the GBM tumor data, we

additionally study shape variability of tumor outlines extracted from single MRI slices, which form planar closed curves. Shapes of white matter tracts in DT-MRI provide information about connectivity of different brain areas. We cluster particular sets of tracts to elucidate connection pathways between Broca's and Wernicke's areas, which are associated with the language circuit. Finally, we use shape models to simulate 3D endometrial tissue shapes, and to define classifiers for ADHD based on shapes of subcortical structures.

Acknowledgements We thank Arvind Rao for sharing the GBM tumor dataset, and acknowledge Joonsang Lee, Juan Martinez, Shivali Narang and Ganesh Rao for their roles in processing the MRIs used to produce the tumor outlines. We thank Zhaohua Ding for providing the DT-MRI tract dataset used in Sect. 24.4.3 and the dataset of subcortical structures for ADHD classification. Finally, we thank Chafik Samir for providing the endometrial tissue surfaces. Shariq Mohammed would like to acknowledge Institutional Research Support from The University of Michigan. Additionally, this work was supported in part by grants NSF DMS-1613054, NSF CCF-1740761, NSF CCF-1839252 and NIH R37-CA214955.

References

1. Bakas, S., Zeng, K., Sotiras, A., Rathore, S., Akbari, H., Gaonkar, B., Rozycki, M., Pati, S., Davatzikos, C.: GLISTRboost: combining multimodal MRI segmentation, registration, and biophysical tumor growth modeling with gradient boosting machines for glioma segmentation. In: *BrainLes*, pp. 144–155 (2015)
2. Bauer, M., Bruveris, M., Harms, P., Møller-Andersen, J.: A numerical framework for Sobolev metrics on the space of curves. *SIAM J. Imag. Sci.* **10**(1), 47–73 (2017)
3. Bharath, K., Kurtek, S., Rao, A., Baladandayuthapani, V.: Radiologic image-based statistical shape analysis of brain tumours. *J. R. Stat. Soc. Ser. C* **67**(5), 1357–1378 (2018)
4. Bhattacharyya, A.: On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.* **35**, 99–109 (1943)
5. Boussejot, R., Kreiseler, D., Schnabel, A.: Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das internet. *Biomedizinische Technik* **40**(1), S317–S318 (1995)
6. Brosens, I., Puttemans, P., Campo, R., Gordts, S., Kinkel, K.: Diagnosis of endometriosis: Pelvic endoscopy and imaging techniques. *Best Pract. Res. Clin. Obstet. Gynaecol.* **18**(2), 285–303 (2004)
7. Čencov, N.N.: *Statistical Decision Rules and Optimal Inference*. American Mathematical Society, Providence (1982)
8. Clifford, G.D., Azuaje, F., McSharry, P.: *Advanced Methods And Tools for ECG Data Analysis*. Artech House, Inc., Boston (2006)
9. De Sousa, F., Melo, E., Vermeulen, L., Fessler, E., Medema, J.P.: Cancer heterogeneity—a multifaceted view. *EMBO Rep.* **14**(8), 686–695 (2013)
10. Dryden, I.L., Mardia, K.V.: *Statistical Shape Analysis*. Wiley, New York (1998)
11. Garg, A., Appel-Cresswell, S., Popuri, K., McKeown, M.J., Beg, M.F.: Morphological alterations in the caudate, putamen, pallidum, and thalamus in Parkinson's disease. *Front. Neurosci.* **9**, 101 (2015)
12. Goldberger, A., Amaral, L., Glass, L., Hausdorff, J., Mark, R., Mietus, J., Moody, G., Peng, C., Stanley, H.: PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation* **101**(23), e215–e220 (2000)

13. Goldsmith, J., Crainiceanu, C.M., Caffo, B.S., Reich, D.S.: Penalized functional regression analysis of white-matter tract profiles in multiple sclerosis. *NeuroImage* **57**(2), 431–439 (2011)
14. Goldsmith, J., Scheipl, F., Huang, L., Wrobel, J., Gellar, J., Harezlak, J., Mclean, M., Swihart, B., Crainiceanu, L.X.C., Reiss, P., Chen, Y., Greven, S., Huo, L., Kunda, M.G., Park, S.Y., Miller, D., Staicu, A.M.: *Refund: Regression with Functional Data* (2018). <https://cran.r-project.org/web/packages/refund/>
15. Holland, E.C.: Glioblastoma multiforme: the terminator. *Proc. Natl. Acad. Sci.* **97**(12), 6242–6244 (2000)
16. Ismail, M., Keynton, R., Mostapha, M., Eltanboly, A., Casanova, M., Gimel'farb, G., El-Baz, A.: Studying autism spectrum disorder with structural and diffusion magnetic resonance imaging: A survey. *Front. Hum. Neurosci.* **10** (2016)
17. Jacobson, S., Jacobson, J., Sokol, R., Martier, S., Chiodo, L.: New evidence for neurobehavioral effects of in utero cocaine exposure. *J. Pediatr.* **129**(4) pp. 581–590 (1996)
18. Jacobson, S., Jacobson, J., Sokol, R., Chiodo, L., Corobana, R.: Maternal age, alcohol abuse history, and quality of parenting as moderators of the effects of prenatal alcohol exposure on 7.5-year intellectual function. *Alcohol. Clin. Exp. Res.* **28**, 1732–1745 (2004)
19. Jermyn, I.H., Kurtek, S., Klassen, E., Srivastava, A.: Elastic shape matching of parameterized surfaces using square root normal fields. In: *European Conference on Computer Vision*, pp. 804–817 (2012)
20. Jermyn, I.H., Kurtek, S., Laga, H., Srivastava, A.: *Elastic Shape Analysis of Three-Dimensional Objects*. Morgan & Claypool Publishers, San Rafael (2017)
21. Joshi, S., Klassen, E., Srivastava, A., Jermyn, I.: Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves. In: *EMMCVPR*, pp. 387–398 (2007)
22. Joshi, S., Xie, Q., Kurtek, S., Srivastava, A., Laga, H.: Surface shape morphometry for hippocampal modeling in Alzheimer's disease. In: *International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pp. 1–8 (2016)
23. Just, N.: Improving tumour heterogeneity MRI assessment with histograms. *Br. J. Cancer* **111**(12), 2205 (2014)
24. Karcher, H.: Riemannian center of mass and mollifier smoothing. *Commun. Pure Appl. Math.* **30**(5), 509–541 (1977)
25. Kass, R.E., Vos, P.W.: *Geometrical Foundations of Asymptotic Inference*, vol. 908. Wiley, New York (2011)
26. Kendall, D.G.: Shape manifolds, procrustean metrics and complex projective spaces. *Bull. Lond. Math. Soc.* **16**, 81–121 (1984)
27. Kurtek, S.: A geometric approach to pairwise Bayesian alignment of functional data using importance sampling. *Electron. J. Stat.* **11**(1), 502–531 (2017)
28. Kurtek, S., Needham, T.: Simplifying transforms for general elastic metrics on the space of plane curves. *ArXiv:1803.10894v1* (2018)
29. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J., Avison, M., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. *IEEE Trans. Med. Imaging* **30**(3), 849–858 (2011)
30. Kurtek, S., Srivastava, A., Klassen, E., Ding, Z.: Statistical modeling of curves using shapes and related features. *J. Am. Stat. Assoc.* **107**(499), 1152–1165 (2012)
31. Kurtek, S., Su, J., Grimm, C., Vaughan, M., Sowell, R., Srivastava, A.: Statistical analysis of manual segmentations of structures in medical images. *Comput. Vis. Image Underst.* **117**, 1036–1050 (2013)
32. Kurtek, S., Wu, W., Christensen, G., Srivastava, A.: Segmentation, alignment and statistical analysis of biosignals with application to disease classification. *J. Appl. Stat.* **40**(6), 1270–1288 (2013)
33. Kurtek, S., Samir, C., Ouchchane, L.: Statistical shape model of elastic endometrial tissue surfaces. In: *International Conference on Pattern Recognition Applications and Methods* (2014)

34. Kurtek, S., Xie, Q., Samir, C., Canis, M.: Statistical model for simulation of deformable elastic endometrial tissue shapes. *Neurocomputing* **173**(P1), 36–41 (2016)
35. Laga, H., Xie, Q., Jermyn, I.H., Srivastava, A.: Numerical inversion of SRNF maps for elastic shape analysis of genus-zero surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(12), 2451–2464 (2017)
36. Le, H.: Locating Frechet means with application to shape spaces. *Adv. Appl. Probab.* **33**(2), 324–338 (2001)
37. Maier-Hein, K.H., Neher, P.F., Houde, J.C., Côté, M.A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.C., Lin, Y.C., Ji, Q., Reddick, W.E., Glass, J.O., Chen, D.Q., Feng, Y., Gao, C., Wu, Y., Ma, J., Renjie, H., Li, Q., Westin, C.F., Deslauriers-Gauthier, S., González, J.O.O., Paquette, M., St-Jean, S., Girard, G., Rheault, F., Sidhu, J., Tax, C.M.W., Guo, F., Mesri, H.Y., Dávid, S., Froeling, M., Heemskerk, A.M., Leemans, A., Boré, A., Pinsard, B., Bedetti, C., Desrosiers, M., Brambati, S., Doyon, J., Sarica, A., Vasta, R., Cerasa, A., Quattrone, A., Yeatman, J., Khan, A.R., Hodges, W., Alexander, S., Romascano, D., Barakovic, M., Auria, A., Esteban, O., Lemkaddem, A., Thiran, J.P., Cetingul, H.E., Odry, B.L., Mailhe, B., Nadar, M.S., Pizzagalli, F., Prasad, G., Villalon-Reina, J.E., Galvis, J., Thompson, P.M., Requejo, F.D.S., Laguna, P.L., Lacerda, L.M., Barrett, R., Dell’Acqua, F., Catani, M., Petit, L., Caruyer, E., Daducci, A., Dyrby, T.B., Holland-Letz, T., Hilgetag, C.C., Stieltjes, B., Descoteaux, M.: The challenge of mapping the human connectome based on diffusion tractography. *Nat. Commun.* **8**(1), 1349 (2017)
38. Marron, J., Ramsay, J.O., Sangalli, L.M., Srivastava, A.: Functional data analysis of amplitude and phase variation. *Stat. Sci.* **30**(4), 468–484 (2015)
39. Oishi, K., Mielke, M., Albert, M., Lyketsos, C., Mori, S.: DTI analyses and clinical applications in Alzheimer’s disease. In: *Advances in Alzheimer’s Disease*, vol. 2, pp. 525–534 (2011)
40. Pennec, X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *J. Math. Imaging Vis.* **25**(1), 127–154 (2006)
41. Ramsay, J., Silverman, B.: *Functional Data Analysis*. Springer, Berlin (2005)
42. Rao, C.R.: Information and the accuracy attainable in the estimation of statistical parameters. In: *Breakthroughs in Statistics*, pp. 235–247. Springer, Berlin (1992)
43. Robinson, D.T.: *Functional Data Analysis and Partial Shape Matching in the Square Root Velocity Framework*. Ph.D. thesis, Florida State University, Tallahassee (2012)
44. Saha, A., Banerjee, S., Kurtek, S., Narang, S., Lee, J., Rao, G., Martinez, J., Bharath, K., Rao, A., Baladandayuthapani, V.: DEMARCATE: Density-based magnetic resonance image clustering for assessing tumor heterogeneity in cancer. *NeuroImage Clin.* **12**, 132–143 (2016)
45. Samir, C., Kurtek, S., Srivastava, A., Canis, M.: Elastic shape analysis of cylindrical surfaces for 3D/2D registration in endometrial tissue characterization. *IEEE Trans. Med. Imaging* **33**(5), 1035–1043 (2014)
46. Srivastava, A., Jermyn, I.H., Joshi, S.H.: Riemannian analysis of probability density functions with applications in vision. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8 (2007)
47. Srivastava, A., Klassen, E., Joshi, S.H., Jermyn, I.H.: Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**, 1415–1428 (2011)
48. Srivastava, A., Wu, W., Kurtek, S., Klassen, E., Marron, J.S.: Registration of functional data using Fisher-Rao metric. *ArXiv:1103.3817v2* (2011)
49. Srivastava, A., Klassen, E.P.: *Functional and Shape Data Analysis*. Springer, Berlin (2016)
50. Wang, L., Beg, M.F., Ratnanather, J.T., Ceritoglu, C., Younes, L., Morris, J.C., Csernansky, J.G., Miller, M.I.: Large deformation diffeomorphism and momentum based hippocampal shape discrimination in dementia of the Alzheimer type. *IEEE Trans. Med. Imaging* **26**(4), 462–470 (2007)
51. Xie, Q., Kurtek, S., Le, H., Srivastava, A.: Parallel transport of deformations in shape space of elastic surfaces. In: *International Conference on Computer Vision* (2013)
52. Yavariabdi, A., Samir, C., Bartoli, A., Ines, D.D., Bourdel, N.: Contour-based TVUS-MR image registration for mapping small endometrial implants. In: *Abdominal Imaging*, vol. 8198, pp. 145–154 (2013)

53. Younes, L.: Computable elastic distance between shapes. *SIAM J. Appl. Math.* **58**(2), 565–586 (1998)
54. Younes, L.: Elastic distance between curves under the metamorphosis viewpoint. [ArXiv:1804.10155](https://arxiv.org/abs/1804.10155) (2018)