

Chapter 12

Elastic Shape Analysis of Surfaces and Images

Sebastian Kurtek, Ian H. Jermyn, Qian Xie, Eric Klassen, and Hamid Laga

Abstract We describe two Riemannian frameworks for statistical shape analysis of parameterized surfaces. These methods provide tools for registration, comparison, deformation, averaging, statistical modeling, and random sampling of surface shapes. A crucial property of both of these frameworks is that they are invariant to reparameterizations of surfaces. Thus, they result in natural shape comparisons and statistics. The first method we describe is based on a special representation of surfaces termed square-root functions (SRFs). The pullback of the \mathbb{L}^2 metric from the SRF space results in the Riemannian metric on the space of surfaces. The second method is based on the elastic surface metric. We show that a restriction of this metric, which we call the partial elastic metric, becomes the standard \mathbb{L}^2 metric under the square-root normal field (SRNF) representation. We show the advantages of these methods by computing geodesic paths between highly articulated surfaces and shape statistics of manually generated surfaces. We also describe applications of this framework to image registration and medical diagnosis.

S. Kurtek (✉)

Department of Statistics, The Ohio State University, Columbus, OH, USA
e-mail: kurtek.1@stat.osu.edu

I.H. Jermyn

Department of Mathematical Sciences, Durham University, Durham, UK
e-mail: i.h.jermyn@durham.ac.uk

Q. Xie

Department of Statistics, Florida State University, Tallahassee, FL, USA
e-mail: qxie@stat.fsu.edu

E. Klassen

Department of Mathematics, Florida State University, Tallahassee, FL, USA
e-mail: klassen@math.fsu.edu

H. Laga

Phenomics and Bioinformatics Research Centre, University of South Australia,
Adelaide, SA, Australia
e-mail: Hamid.Laga@unisa.edu.au

12.1 Introduction

Shape is an important physical property of 3D objects that helps characterize their appearance. As a result, statistical shape analysis, which is concerned with quantifying shape as a random object and developing tools for generating shape registrations, comparisons, deformations, averages, probability models, hypothesis tests, Bayesian estimates, and other statistical procedures on shape spaces, plays an important role in many applications including medical imaging, biometrics, bioinformatics, 3D printing, and computer graphics. Medical imaging is a primary example of an application where shape statistics can play a very important role. Advances in noninvasive imaging technology, such as magnetic resonance imaging (MRI), have enabled researchers to study biological variations of anatomical structures. Studying shapes of 3D anatomies is of particular interest because many complex diseases can potentially be linked to alterations of these shapes. Thus, statistical shape analysis can be a central tool in disease diagnosis and design of novel treatment strategies. The two methods described in this chapter have been successfully applied to classification of attention deficit hyperactivity disorder (ADHD) [31, 32, 47] and mathematics deficiency [35] using shapes of subcortical structures, and statistical modeling of endometrial tissues [41] in the presence of endometriosis.

In this chapter, we focus on shape analysis of parameterized surfaces. In particular, we describe two recent Riemannian frameworks that allow comparison, matching, deformation, averaging, and statistical modeling of observed shapes. This work was motivated by the widespread success of elastic Riemannian shape analysis of curves [27, 37, 40, 42, 51]. The main benefit of Riemannian shape analysis is in the breadth of mathematical tools at our disposal, resulting in a *principled* statistical shape analysis framework. We note that there are currently very few Riemannian approaches to shape analysis of 3D objects. Similar to curves, researchers have proposed many different representations of surfaces. Many groups study shapes of surfaces by embedding them in volumes and deforming these volumes under the large deformation diffeomorphic metric mapping (LDDMM) framework [11, 12, 20, 26, 46]. While these methods are both prominent and pioneering in medical image analysis, they are typically computationally expensive since they try to match not only the objects of interest but also some background space containing them. An important benefit of the LDDMM framework is that it utilizes the Riemannian geometry of the reparameterization group to compute shape comparisons and deformations. A closely related approach utilizes inner metrics to describe shape deformations, which are prescribed directly on the surface [3]. Others study 3D shapes using manually generated landmarks under Kendall's shape theory [15], level sets [39], curvature flows [22], point clouds [2], or medial axis representations [5, 19].

However, the most natural representation for studying the shape of a 3D object would seem to be a parameterized surface. In this case, there is an additional difficulty in handling the parameterization variability. Specifically, a reparameterization of a surface (achieved using an appropriate function $\gamma \in \Gamma$ made

precise later) does not change the shape of the object. Thus, a main goal in shape analysis is to define Riemannian metrics and subsequent statistical analyses, which are invariant to the introduction of an arbitrary parameterization in shape representations. Methods such as SPHARM [6, 29] or SPHARM-PDM [17, 44] tackle this problem by choosing a fixed parameterization that is analogous to the arc-length parameterization on curves. Kilian et al. [30] presented a technique for computing geodesics between triangulated meshes (discretized surfaces) but at their given parameterizations, thus requiring the registration problem to be solved manually or using some other available method. In fact, a large set of papers in the literature treat surface registration as a preprocessing step [28]. In such methods, points across surfaces are first registered using some predefined energy functions such as the entropy [8] or the minimum description length [13]. Once the surfaces are registered, they are compared using standard procedures. There are several fundamental problems with such approaches; first, the energy used for registration does not lead to a proper distance on the shape space of surfaces. Second, the registration procedure is typically completely unrelated to the rest of the analysis. In other words, the two tasks are performed under different metrics. Figure 12.1 displays the various representations of surfaces used for shape analysis.

The remainder of this chapter describes two Riemannian frameworks for statistical shape analysis of parameterized surfaces that overcome the above presented difficulties. In particular, the defined Riemannian metrics are invariant to reparameterizations of surfaces and allow shape comparisons via geodesic paths in the shape space. Geodesics can in turn be used to define statistics on shape spaces including the Karcher mean and covariance. Tools for other statistical procedures are also presented including principal component analysis and random sampling from

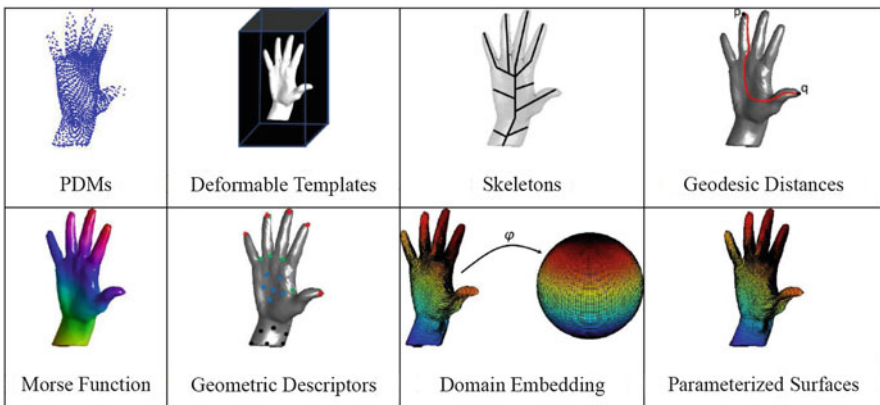


Fig. 12.1 Different representations of surfaces for shape analysis. (courtesy of Srivastava et al. [43])

Gaussian-type shape models. Finally, an application to classification of attention deficit hyperactivity disorder is presented, where the benefits of these methods are seen through a superior classification rate.

12.2 Surface Representations and Riemannian Metrics

Let \mathcal{F} be the space of all smooth embeddings $f : D \rightarrow \mathbb{R}^3$, where D represents the surface domain. Depending on the application of interest D can be the sphere for closed surfaces, disk for hemispherical surfaces, or square for quadrilateral surfaces. Each such embedding defines a surface $f(D) \subset \mathbb{R}^3$. Let Γ be the set of all diffeomorphisms of D . Γ will act as the reparameterization group for surfaces. It acts naturally on \mathcal{F} by composition: $(f, \gamma) \mapsto f \circ \gamma$. Thus, the space of surfaces can be thought of as the quotient space \mathcal{F}/Γ , i.e. the space of equivalence classes under the action of Γ endowed with appropriate structures.

Since \mathcal{F} is a vector space, the tangent space at any point $f \in \mathcal{F}$ denoted by $T_f(\mathcal{F})$ can be identified with \mathcal{F} . Given the Lebesgue measure ds on D , one can define a Riemannian structure on \mathcal{F} as follows. For $\delta f_1, \delta f_2 \in T_f(\mathcal{F})$, define the \mathbb{L}^2 Riemannian metric as $\langle \delta f_1, \delta f_2 \rangle = \int_D \langle \delta f_1(s), \delta f_2(s) \rangle ds$, where the inner product inside the integral is the Euclidean inner product in \mathbb{R}^3 . The resulting squared \mathbb{L}^2 distance between two surfaces $f_1, f_2 \in \mathcal{F}$ is $\int_D |f_1(s) - f_2(s)|^2 ds$. In this expression, $|\cdot|$ denotes the standard two-norm of a vector in \mathbb{R}^3 . While simple, this metric has a critical defect: just as in the case of curves, the action of Γ does not preserve distances. In other words, the group Γ does not act on the space \mathcal{F} by isometries under the \mathbb{L}^2 metric. This is easily seen through the following expression: $d(f_1 \circ \gamma, f_2 \circ \gamma)^2 = \int_D |f_1(\gamma(s)) - f_2(\gamma(s))|^2 ds = \int_D |f_1(\tilde{s}) - f_2(\tilde{s})|^2 J_\gamma(s)^{-1} d\tilde{s} \neq d(f_1, f_2)^2$, where J_γ is the determinant of the Jacobian of γ . In this equation, we have used the substitution $\tilde{s} = \gamma(s)$ and $J_\gamma(s) ds = d\tilde{s}$. The inequality comes from the fact that, in general, γ is not area preserving and thus the determinant of the Jacobian is not one at all points. The lack of isometry means that the shape space \mathcal{F}/Γ does not inherit the structure of a Riemannian manifold from \mathcal{F} , thereby making this metric difficult to use for analyzing shapes of parameterized surfaces. One solution is to restrict attention to area-preserving diffeomorphisms [21], but this restriction proves to be very limiting in practice. Another solution is to develop a new representation of surfaces such that the action of Γ preserves \mathbb{L}^2 distances. Then, one can use the pullback of the \mathbb{L}^2 metric from the new representation space to define a Riemannian structure on the space \mathcal{F} . We take this approach and present two different representations of surfaces that satisfy these conditions: the square-root function (SRF) and the square-root normal field (SRNF).

12.2.1 Square-Root Function Representation and Pullback Metric

Let $(x, y) : D \rightarrow \mathbb{R}^2$ be coordinates on D ; then, $f_x(s) = \frac{\partial f}{\partial x}(s)$ and $f_y(s) = \frac{\partial f}{\partial y}(s)$. To endow \mathcal{F} with a Riemannian metric, we begin by defining a new representation of surfaces called square-root functions or SRFs [32, 33, 38]:

Definition 1. Define the mapping $Q : \mathcal{F} \rightarrow \mathbb{L}^2$ as $Q(f)(s) = q(s) = \sqrt{|n(s)|}f(s)$, where $n(s) = f_x(s) \times f_y(s)$ is the normal vector to the surface f at point s .

The factor $|n(s)|$ can be interpreted as the ratio of infinitesimal areas of the surface at $f(s)$ and the domain at s , the “area multiplication factor.” For any $f \in \mathcal{F}$, we will refer to $q(s) = Q(f)(s)$ as the SRF of f . Since we defined \mathcal{F} as the space of smooth surfaces, $Q(\mathcal{F})$ is a subset of $\mathbb{L}^2(D, \mathbb{R}^3)$, henceforth denoted \mathbb{L}^2 . If a surface f is reparameterized by γ , then its SRF changes to $(q, \gamma) = (q \circ \gamma)\sqrt{J_\gamma}$. This can be extended to a right action of Γ on \mathbb{L}^2 .

We choose the natural \mathbb{L}^2 metric on the space of SRFs: the inner product of any two elements $\delta q_1, \delta q_2 \in T_q(\mathbb{L}^2)$ is $\langle \delta q_1, \delta q_2 \rangle = \int_D \langle \delta q_1(s), \delta q_2(s) \rangle ds$. This metric has the key property that Γ acts by isometries on \mathbb{L}^2 . As a result, if we pullback this metric to \mathcal{F} , the resulting Riemannian metric is also preserved by the action of Γ , unlike the plain \mathbb{L}^2 metric on \mathcal{F} mentioned at the beginning of this section. To obtain the pullback metric, we must first derive the differential of the mapping Q at any surface f , denoted by $Q_{*,f}$. This is a linear mapping between tangent spaces $T_f(\mathcal{F})$ and $T_{Q(f)}(\mathbb{L}^2)$. For a tangent vector $\delta f \in T_f(\mathcal{F})$, the mapping $Q_{*,f} : T_f(\mathcal{F}) \rightarrow T_{Q(f)}(\mathbb{L}^2)$ is given by

$$Q_{*,f}(\delta f) = \frac{1}{2|n|^{\frac{3}{2}}}(n \cdot n_{\delta f})f + \sqrt{|n|} \delta f. \quad (12.1)$$

The quantity $n_{\delta f}$ depends on both f and δf and is defined as $f_x \times \delta f_y + \delta f_x \times f_y$. The pullback metric on \mathcal{F} is then defined as usual.

Definition 2. For any $f \in \mathcal{F}$ and any $\delta f_1, \delta f_2 \in T_f(\mathcal{F})$, define the inner product

$$\langle \langle \delta f_1, \delta f_2 \rangle \rangle_f = \langle Q_{*,f}(\delta f_1), Q_{*,f}(\delta f_2) \rangle, \quad (12.2)$$

where the inner product on the right side is the standard inner product in \mathbb{L}^2 .

To write the metric in Definition 2 in full detail, we use the expression for $Q_{*,f}(\delta f)$ given in Eq. (12.1):

$$\begin{aligned} \langle \langle \delta f_1, \delta f_2 \rangle \rangle_f &= \left\langle \frac{1}{2|n|^{\frac{3}{2}}}(n \cdot n_{\delta f_1})f + \sqrt{|n|} \delta f_1, \frac{1}{2|n|^{\frac{3}{2}}}(n \cdot n_{\delta f_2})f + \sqrt{|n|} \delta f_2 \right\rangle \\ &= \left\langle \frac{1}{4|n|^3}(n \cdot n_{\delta f_1})f, (n \cdot n_{\delta f_2})f \right\rangle + \left\langle \frac{1}{2|n|}[(n \cdot n_{\delta f_2})\delta f_1 + (n \cdot n_{\delta f_1})\delta f_2], f \right\rangle + \langle |n|\delta f_1, \delta f_2 \rangle. \end{aligned}$$

As stated, because of the structure of Q , the action of Γ on \mathcal{F} is by isometries under this metric. That is, for any surface $f \in \mathcal{F}$, a $\gamma \in \Gamma$, and two tangent vectors $\delta f_1, \delta f_2 \in T_f(\mathcal{F})$, we have $\langle\langle \delta f_1 \circ \gamma, \delta f_2 \circ \gamma \rangle\rangle_{f \circ \gamma} = \langle\langle \delta f_1, \delta f_2 \rangle\rangle_f$.

It is frequently important for shape analysis to be invariant to changes in the positions, orientations, and sizes of the surfaces being compared. In other words, the metric should be invariant to the action of the (direct) similarity group of translations, rotations, and scalings. One way to achieve such invariance is by normalizing, i.e. by picking a distinguished representative of each equivalence class under the group action, and then computing distances only between these distinguished elements. Thus, translations may be removed by centering, $f_{\text{centered}}(s) = f(s) - \frac{\int_D f(s)|n(s)|ds}{\int_D |n(s)|ds}$, while scalings may be removed by rescaling all surfaces to have unit area, $f_{\text{scaled}}(s) = \frac{f(s)}{\sqrt{\int_D |n(s)|ds}}$. Slightly abusing notation, we denote the space of such normalized surfaces also by \mathcal{F} . Then, \mathcal{F} forms the ‘‘preshape space’’ in our analysis. Paired with the Riemannian metric $\langle\langle \cdot, \cdot \rangle\rangle$, it becomes a Riemannian manifold.

Another, often equivalent, way to achieve invariance is by constructing the quotient space under the group action, inducing a metric on the quotient space from the covering space and then computing distances between points in the quotient space (i.e. between equivalence classes). This is how we will deal with the actions of the rotation group $SO(3)$ and Γ . The rotation group acts on \mathcal{F} according to $(O, f) = Of$, for $O \in SO(3)$ and $f \in \mathcal{F}$. It is easy to check that this action is by isometries. We have already seen that the action of Γ is by isometries too. Furthermore, the actions of Γ and $SO(3)$ on \mathcal{F} commute, allowing us to define an action of the product group. The equivalence class or orbit of a surface f is given by $[f] = \{O(f \circ \gamma) | O \in SO(3), \gamma \in \Gamma\}$, and the set of all $[f]$ is by definition the quotient space $\mathcal{S} = \mathcal{F}/\Gamma = \{[f] | f \in \mathcal{F}\}$. This quotient space is called the ‘‘shape space.’’

The next step is to define geodesic paths and distances in the shape space \mathcal{S} . This is accomplished using the following joint optimization problem. Let f_1 and f_2 denote two surfaces and let $\langle\langle \cdot, \cdot \rangle\rangle$ be the Riemannian metric on \mathcal{F} . Then, the geodesic distance between shapes of f_1 and f_2 is given by

$$d([f_1], [f_2]) = \min_{(O, \gamma) \in SO(3) \times \Gamma} \left(\min_{F : [0, 1] \rightarrow \mathcal{F}} \left(\int_0^1 \langle\langle \frac{dF(t)}{dt}, \frac{dF(t)}{dt} \rangle\rangle^{(1/2)} dt \right) \right) \tag{12.3}$$

(12.3)

In this equation, $F(t)$ is a path in \mathcal{F} indexed by t . The quantity $L(F) = \int_0^1 \langle\langle \frac{dF(t)}{dt}, \frac{dF(t)}{dt} \rangle\rangle^{(1/2)} dt$ provides the length of the path F . The minimization inside the brackets represents the problem of finding a geodesic path between the surfaces f_1 and $O(f_2 \circ \gamma)$, where O and γ stand for an arbitrary rotation and reparameterization of f_2 , respectively. This is computed using a path-straightening technique. We omit the details of this method here and refer the interested reader to [34] for

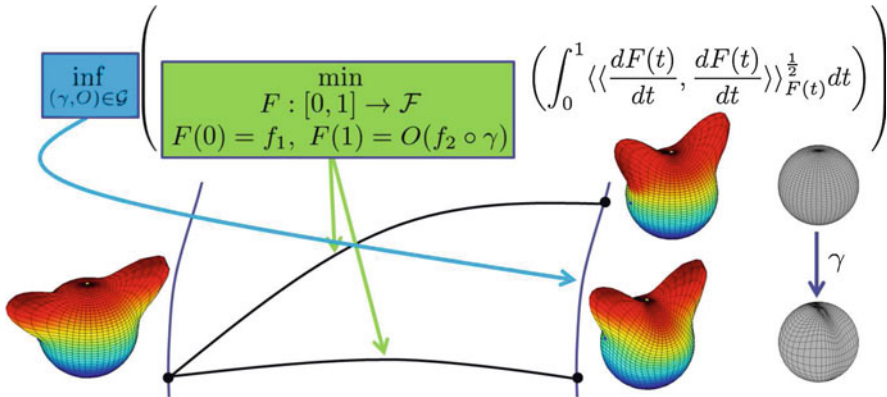


Fig. 12.2 Pictorial description of the process of computing geodesics in the shape space of surfaces

details. The minimization outside the brackets seeks the optimal rotation and reparameterization of the second surface so as to best match it with the first surface. This optimization is performed iteratively using Procrustes analysis to solve for optimal rotations and a gradient descent algorithm [32, 33] to solve for the optimal reparameterization. A few details of this registration procedure are presented in the next section. In simple words, the outside optimization solves the registration problem while the inside optimization solves for both an optimal deformation (geodesic path, F^*) and a formal geodesic distance between shapes. Figure 12.2 displays the joint optimization problem defined in Eq. (12.3). Figure 12.3 displays several examples of geodesic comparisons for complex surfaces with many articulated parts. We note the clear benefit of finding optimal reparameterizations during the geodesic computation. The geodesics in the shape space are much more natural than those in the pre-shape space. Furthermore, the decrease in the distance due to optimization over the reparameterization group is significant in all of the presented examples. These properties will additionally lead to improved shape statistics and more parsimonious shape models.

12.2.2 Application to Image Registration

The problem of image registration is common across multiple application areas. Given a set of observed images, the goal is to establish point correspondence across the domains of these images. Although the registration problem has been studied for almost two decades, there continue to be some fundamental limitations in the popular solutions that make them suboptimal, difficult to evaluate, and limited in scope. To explain these limitations, let f_1 and f_2 represent two \mathbb{R}^n -valued images. A pairwise registration between these images is defined as finding a diffeomorphic

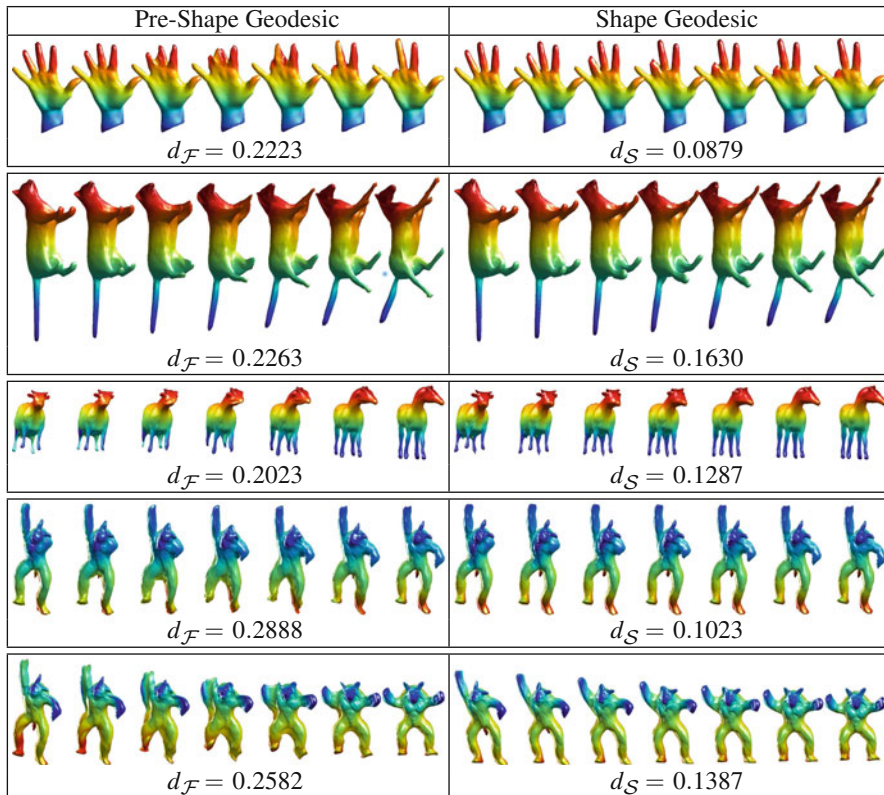


Fig. 12.3 Comparison of geodesics computed under $\langle\langle \cdot, \cdot \rangle\rangle$ in the pre-shape space and shape space

mapping γ of the image domain D to itself, such that pixels $f_1(s)$ and $f_2(\gamma(s))$ are optimally matched to each other for all $s \in D$. To develop an algorithm for registration one needs an objective function for formalizing the notion of optimality. A common type of objective function used in registration frameworks is $\mathcal{L}(f_1, f_2 \circ \gamma) = \int_D |f_1(s) - f_2(\gamma(s))|^2 ds + \lambda R(\gamma)$, where R is a regularization penalty on γ and λ is a positive constant. This objective function has many shortcomings including lack of symmetry (registration of f_1 to f_2 is different from registration of f_2 to f_1). Next, we outline several important properties for image registration and show that the framework for shape analysis of surfaces can be applied to the image registration problem. Furthermore, this framework overcomes many of the shortcomings of current registration methods.

In the case of images we extend the definition of \mathcal{F} to all images on some domain D that take value in \mathbb{R}^n , $n > 1$. (When dealing with grayscale images, we add the gradient of the image f to create an image in \mathbb{R}^3 .) We let Γ denote the reparameterization group, also called the image-warping group. Let $\mathcal{L}(f_1, (f_2, \gamma))$ denote the objective function for matching f_1 and f_2 by optimizing over Γ (here γ is

assumed to be applied to f_2 resulting in $(f_2, \gamma) \in \mathcal{F}$. Then, the desired properties of \mathcal{L} are (for any $f_1, f_2 \in \mathcal{F}$ and $\gamma \in \Gamma$): (1) symmetry; (2) positive definiteness; (3) lack of bias: if f_1, f_2 are constant functions then $\mathcal{L}(f_1, f_2) = \mathcal{L}(f_1, (f_2, \gamma))$; (4) invariance to identical warping: $\mathcal{L}(f_1, f_2) = \mathcal{L}((f_1, \gamma), (f_2, \gamma))$; (5) triangle inequality; and (6) Γ is a group with composition as the group action. These properties have been discussed previously in [9, 45, 49].

Next, we define a representation of images, similar to the SRF, which allows invariance to Γ under the \mathbb{L}^2 metric [48, 49].

Definition 3. Define the mapping $\tilde{Q} : \mathcal{F} \rightarrow \mathbb{L}^2(D, \mathbb{R}^n)$, $n > 1$, as $\tilde{Q}(f)(s) = \tilde{q}(s) = \sqrt{|a(s)|}f(s)$, where $|a(s)| = |f_x(s) \wedge f_y(s)|$, where \wedge is the wedge product.

We will refer to this representation as the extended SRF. The extended SRF is simply a generalization of the SRF used for surfaces to functions taking values in \mathbb{R}^n , $n > 1$. Assuming the original set of images to be smooth, the set of all extended SRFs is a subset of \mathbb{L}^2 . One can show that the action of Γ on \mathbb{L}^2 is exactly the same, *mutatis mutandis*, as that in the previous section. This implies that this group acts on \mathbb{L}^2 by isometries, satisfying property (4). This leads to the following definition of the objective function.

Definition 4. Define an objective function for registration of any two images f_1 and f_2 , represented by their extended SRFs \tilde{q}_1 and \tilde{q}_2 , as $\mathcal{L}(f_1, (f_2, \gamma)) = \|\tilde{q}_1 - (\tilde{q}_2, \gamma)\|$.

The registration is then achieved by minimizing this objective function: $\gamma^* = \operatorname{arginf}_{\gamma \in \Gamma} \mathcal{L}(f_1, (f_2, \gamma)) = \operatorname{arginf}_{\gamma \in \Gamma} \|\tilde{q}_1 - (\tilde{q}_2, \gamma)\|$. The objective function \mathcal{L} given in Definition 4 satisfies all of the properties listed earlier. The \mathbb{L}^2 norm between extended SRFs of images becomes a proper measure of registration between images since it remains the same if the pixel correspondence is unchanged. This leads to a quantity that serves as both the registration objective function and an extrinsic distance between registered images ($\|\tilde{q}_1 - (\tilde{q}_2, \gamma^*)\|$). It is important to note that the proposed objective function has only one term (similarity term) and the regularity term appears to be missing. However, the similarity term has built-in regularity, since it includes the determinant of the Jacobian of the transformation γ . Additional regularity can also be introduced as in the LDDMM framework [4].

Gradient Descent Method for Optimization Over Γ The optimization problem over Γ is a major component of this registration framework and we use a gradient descent method to solve it. Since Γ is a group, we use the gradient to solve for the incremental warping γ , on top of the previous cumulative warping γ_o , as follows. First, define a cost function with respect to γ as $E(\gamma) = \|\tilde{q}_1 - \phi_{\tilde{q}_2}(\gamma)\|^2$, where $\phi_{\tilde{q}} : \Gamma \rightarrow [\tilde{q}]$ is defined to be $\phi_{\tilde{q}}(\gamma) = (\tilde{q}, \gamma)$ and $\tilde{q}_2^o = (\tilde{q}_2, \gamma_o)$ with γ_o being the current registration function. Given a set of orthonormal basis elements, say \mathbf{B} , of $T_{\gamma_{id}}(\Gamma)$, the gradient at γ_{id} takes the form $\nabla E(\gamma_{id}) = \sum_{b \in \mathbf{B}} \langle \tilde{q}_1 - \phi_{\tilde{q}_2}(\gamma_{id}), \phi_{\tilde{q}_2^o, *}(b) \rangle$. In this equation, $\phi_{\tilde{q}_2^o, *}(b)$ denotes the differential of $\phi_{\tilde{q}_2}$ in the direction of b and brackets denote the \mathbb{L}^2 inner product. We omit the derivation of $\phi_{\tilde{q}, *}$ and note that it is the same as presented in [33]. We note that this gradient-based solution can also be employed in the search for optimal γ in the case of parameterized surfaces.

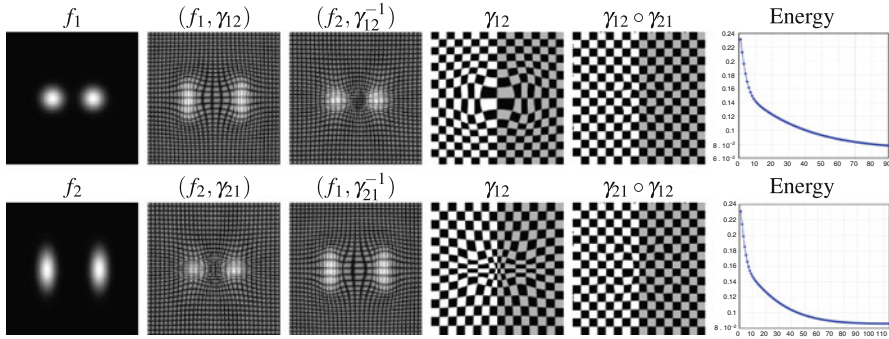


Fig. 12.4 Registration of two synthetic grayscale images. $\gamma_{12} = \operatorname{argmin}_{\gamma \in \Gamma} \mathcal{L}(f_2, (f_1, \gamma))$ and $\gamma_{21} = \operatorname{argmin}_{\gamma \in \Gamma} \mathcal{L}(f_1, (f_2, \gamma))$. $\|q_1 - q_2\| = 0.2312$, $\|\tilde{q}_1 - (\tilde{q}_2, \gamma_{21})\| = 0.0728$, and $\|(\tilde{q}_1, \gamma_{12}) - \tilde{q}_2\| = 0.0859$. (Courtesy of Xie et al. [49])

There, one has to perform an additional search over the rotation group $SO(3)$, and then solve for the geodesic using path straightening.

In Fig. 12.4 we display an example of registering two smooth grayscale images using this framework. The correspondence appears to be good, and more importantly, the resulting distance between the registered images is approximately symmetric and the registration is inverse consistent. That is, the compositions $\gamma_{21} \circ \gamma_{12}$ and $\gamma_{12} \circ \gamma_{21}$ result in the identity mapping. This provides empirical support for the claim that this method satisfies all properties outlined earlier.

12.2.3 Elastic Riemannian Metric

While the SRF representation has its advantages, it has two important drawbacks. First, there is no intuitive interpretation or justification for the use of the metric, unlike the elastic Riemannian metric used in the case of curves; rather, it was solely devised for the convenience of being able to compare the shapes of parameterized surfaces using the \mathbb{L}^2 metric in the SRF space. Second, the associated metric is not translation invariant; translating two shapes equally does not preserve the distance between them, which can cause some issues during statistical analysis. To overcome these two drawbacks, while preserving the advantages of the SRF representation, a different representation of surfaces was introduced in [25]: the square-root normal field, or SRNF, the details of which we present next.

We first recall the elastic metric in the case of curves [42]. Let $\beta : D \rightarrow \mathbb{R}^2$ be a curve, where D temporarily denotes either the circle S^1 or the interval. Let $r ds = |\dot{\beta}| ds$ be the induced metric measure, where ds is Lebesgue measure on D , and let \tilde{n} be the normalized normal vector to the curve. We can represent the curve uniquely up to translations by the pair (r, \tilde{n}) . Then, the family of elastic metrics takes the form

$$\langle\langle(\delta r_1, \delta \tilde{n}_1), (\delta r_2, \delta \tilde{n}_2)\rangle\rangle_{(r, \tilde{n})} = 2\lambda \int_D \frac{\delta r_1(s)\delta r_2(s)}{r(s)} ds + c \int_D \langle\delta \tilde{n}_1(s), \delta \tilde{n}_2(s)\rangle r(s) ds, \tag{12.4}$$

where $(\delta r_i, \delta \tilde{n}_i), i \in \{1, 2\}$ are tangent vectors at the curve (r, \tilde{n}) , and $\lambda, c \in \mathbb{R}_+$.

Although this metric was defined for curves, it can immediately be applied to surfaces, with $r = |n|$ denoting the induced metric measure and $\tilde{n} = \frac{n}{r}$ denoting the unit normal vector. We thereby define a new Riemannian metric on the space of parameterized surfaces known, for reasons that will become clear, as the ‘‘partial elastic metric.’’ To understand this metric, consider that a small change in f on an infinitesimal patch in the surface around the point $f(s)$ can be decomposed into a change in the normal direction of the patch (‘‘bending’’) and a change in its geometry (‘‘stretching’’). Since a change in $\tilde{n}(s)$ corresponds to a change in normal direction and a change in r corresponds to a change in the area of the patch, we see that this metric has an interpretation directly analogous to its interpretation in the case of curves.

However, when D is two-dimensional, the change in area does not completely characterize a change in the geometry of a patch. A change in geometry can be decomposed into a change in area and an area-preserving change of shape [10, 14]. The partial elastic metric measures the first type of change, but does not measure changes in f that change the shape of a patch while preserving its area and normal direction. This limitation came about because the correspondence that we used between the quantity r in the case of curves and in the case of surfaces was in fact incomplete. We interpreted r for curves as the induced metric measure, but r^2 is also the full induced metric; in one dimension there is no difference. This suggests that instead of using (r, \tilde{n}) to represent a surface, we should use (g, \tilde{n}) , where $g = f^*h$ is the full pullback metric (with h the metric on \mathbb{R}^3). The metric g contains more information than r , because r is just $|g|^{\frac{1}{2}}$, where $|\cdot|$ indicates the determinant.

This in turn suggests that Eq. (12.4) is merely a special case of a more general elastic metric for surfaces, and indeed this is the case. This ‘‘full elastic metric’’ is defined, up to an overall scale, by

$$\begin{aligned} &\langle\langle\delta g_1, \delta \tilde{n}_1\rangle, \langle\delta g_2, \delta \tilde{n}_2\rangle\rangle_{(g, \tilde{n})} \\ &= \int_D \left[\text{tr}(g^{-1}\delta g_1 g^{-1}\delta g_2) + \frac{\lambda}{2} \text{tr}(g^{-1}\delta g_1) \text{tr}(g^{-1}\delta g_2) + c \langle\delta \tilde{n}_1, \delta \tilde{n}_2\rangle \right] |g|^{\frac{1}{2}} ds, \end{aligned} \tag{12.5}$$

where, for positivity, $\lambda > \frac{-2}{3}$ and $c \in \mathbb{R}_+$.¹

To see that the partial metric is indeed a special case of this full metric, note that the term multiplied by λ can be written as

¹It is interesting to note that the first two terms in the full elastic metric form the unique family of ultralocal metrics on the space of Riemannian metrics on which diffeomorphisms act by isometries [14]. The $\lambda = 0$ case of this metric on Riemannian metrics has been studied in e.g., [10, 16, 18].

$$\int_D ds |g|^{\frac{1}{2}} \text{tr}(g^{-1} \delta g_1) \text{tr}(g^{-1} \delta g_2) = 4 \int_D ds |g|^{-\frac{1}{2}} \delta_1(|g|^{\frac{1}{2}}) \delta_2(|g|^{\frac{1}{2}}). \tag{12.6}$$

Since $r = |g|^{\frac{1}{2}}$, this is the same as the first term in Eq. (12.4), while the last term is the same in each case. Unlike the partial metric, however, the first two terms in Eq. (12.5) measure not only changes in local area but also any “stretching” of the surface, i.e., changes in both the area and the shape of local patches. The third term continues to measure changes in the normal direction, that is, “bending.” The intuitive interpretation of the metric thus remains unchanged, but now all types of change in f are measured. Indeed, the map from \mathcal{F} to the (g, \tilde{n}) representation is bijective, up to translations (although it is not surjective) [1].

Having defined this metric for surfaces, i.e., codimension-1 embedded submanifolds of \mathbb{R}^3 , it is easy to see that it applies in fact to codimension-1 embedded submanifolds in any number of dimensions, and, with a simple generalization of the third term, to embedded submanifolds of any codimension in any number of dimensions. In particular, when D is one-dimensional, and thus g is a scalar, the first two terms become the same, so that the full metric becomes identical to the partial metric (which is thus no longer partial), which is, in turn, just the elastic metric for curves. Equation (12.5) is thus the most general form of the elastic metric and deserves further study. However, we defer further analysis of its properties to another place. We now focus on the partial metric.

12.2.4 Square-Root Normal Field Representation of Surfaces

An important and remarkable property of the partial elastic metric in Eq. (12.4) is that, despite appearances, for a particular choice of the ratio of λ and c , it is Euclidean: i.e., we can find a transformation of the representation such that this metric takes on the simple \mathbb{L}^2 form. This is strictly analogous to the case of the elastic metric for curves, which gave rise to the square-root velocity function (SRVF) [42]. Just as in this case, and in the case of the SRF, the existence of a representation that transforms the metric to \mathbb{L}^2 form means that parts of the shape analysis framework can be dramatically simplified.

This new, convenient representation of surfaces is called the square-root normal field (SRNF) and is defined as follows [25]:

Definition 5. Define the mapping $H : \mathcal{F} \rightarrow \mathbb{L}^2$ as $H(f)(s) = h(s) = \sqrt{r(s)} \tilde{n}(s) = \frac{n(s)}{\sqrt{r(s)}} = \frac{n(s)}{\sqrt{|n(s)|}}$.

This is strictly analogous to the SRVF defined for curves, except that there the tangent vector is used instead of the normal vector/form. Since $|h|^2 = r$, the \mathbb{L}^2 norm of h is just the area of the surface (again, *cf.* the case of curves, where the \mathbb{L}^2 norm of the SRVF gives the length of the curve). Thus, just as in the case of SRFs, the space of SRNFs is also a subset of $\mathbb{L}^2(D, \mathbb{R}^3)$ or simply \mathbb{L}^2 .

We now show that the \mathbb{L}^2 metric in the SRNF space is the same as that in Eq. (12.4). The derivative map between the tangent spaces at (r, \tilde{n}) and $h \in \mathbb{L}^2$ is given by

$$\delta h(s) = \frac{1}{2\sqrt{r(s)}}\tilde{n}(s)\delta r(s) + \sqrt{r(s)}\delta\tilde{n}(s). \quad (12.7)$$

Taking the \mathbb{L}^2 inner product between two such vectors, we obtain

$$\langle \delta h_1, \delta h_2 \rangle_h = \frac{1}{4} \int_D \frac{\delta r_1(s)\delta r_2(s)}{r(s)} ds + \int_D r(s) \langle \delta\tilde{n}_1(s), \delta\tilde{n}_2(s) \rangle ds \quad (12.8)$$

since $\langle \tilde{n}(s), \delta\tilde{n}_i(s) \rangle = 0$. This is just the partial elastic metric for $\lambda = 1/8$ and $c = 1$. We thus find that if $c/\lambda = 8$, the SRNF representation acts as ‘‘Euclidean coordinates’’ for the partial elastic metric, bringing it to \mathbb{L}^2 form.

As in the case of SRFs, we must also remove all shape-preserving transformations in order to generate shape comparisons. A major advantage of the SRNF is that it (and consequently the partial and full elastic metrics) is automatically invariant to translations, simply because it depends only on derivatives of f . As previously, we can scale all surfaces to have unit area. It is easily checked that rotating a surface $f \mapsto Of$ sends $h \mapsto Oh$ and reparameterizing a surface $f \mapsto f \circ \gamma$ sends $h \mapsto (h, \gamma) = (q \circ \gamma)\sqrt{J_\gamma}$. It is also easy to verify that just as in the SRF case, the action of Γ on the space of surfaces under the partial elastic metric is by isometries. This allows the use of this metric for parameterization-invariant shape analysis.

In order to generate comparisons of shapes using geodesic paths and distances, we take a similar approach to that presented in Sect. 12.2.1 for the SRF representation. Because the action of Γ on the SRNF space is the same as that on the SRF space, the gradient descent algorithm to compute optimal reparameterizations remains unchanged. To compute geodesics, we also use the path-straightening approach [41]. In Fig. 12.5, we again provide a comparison of geodesics computed in the pre-shape space and shape space under the partial elastic metric. We use the same examples as in Fig. 12.3. As previously, the geodesics in the shape space have much lower distances due to the additional optimization over Γ . They are also much more natural due to the improved correspondence of geometric features across surfaces. When comparing the shape space results in this figure to those in Fig. 12.3, we notice that the partial elastic metric provides better registrations than the pullback metric from SRF space. This is expected due to the nice stretching and bending interpretation of this metric.

12.3 Shape Statistics

In this section, we present tools for computing two fundamental shape statistics, the mean, and the covariance. We then utilize these quantities to estimate a generative

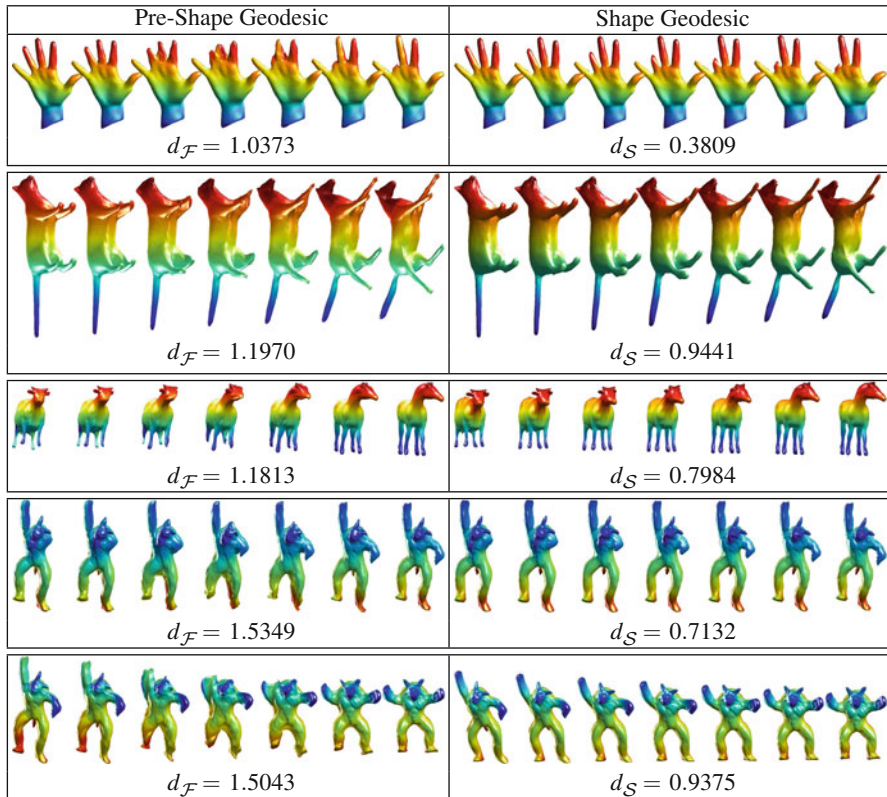


Fig. 12.5 Comparison of geodesics computed under $\langle\langle\cdot, \cdot\rangle\rangle$ in the pre-shape space and shape space

Gaussian model and draw random samples from this model. These methods have been previously presented for SRFs and SRNFs in [31, 36].

First, we define an intrinsic shape mean, called the Karcher mean. Let $\{f_1, f_2, \dots, f_n\} \in \mathcal{F}$ denote a sample of surfaces. Also, let F_i^* denote a geodesic path (in the shape space) between the shape of some surface f and the shape of the i th surface in the sample, f_i . Then, the sample Karcher mean shape is given by

$$[\bar{f}] = \operatorname{argmin}_{[f] \in \mathcal{S}} \sum_{i=1}^n L(F_i^*)^2.$$

A gradient-based approach for finding the Karcher mean is given in, for example, [15] and is omitted here for brevity. Note that the resulting Karcher mean shape is an entire equivalence class of surfaces.

Once the sample Karcher mean has been computed, the evaluation of the Karcher covariance is performed as follows. First, find the shooting vectors from the estimated Karcher mean $\bar{f} \in [\bar{f}]$ to each of the surface shapes in the sample, $v_i = \frac{dF_i^*}{dt}|_{t=0}$, where $i = 1, 2, \dots, n$ and F^* denotes a geodesic path in \mathcal{S} . This is

accomplished using the inverse exponential map, which is used to map points from the representation space to the tangent space. Then, perform principal component analysis by applying the Gram–Schmidt procedure (under the chosen metric $\langle\langle \cdot, \cdot \rangle\rangle$) to generate an orthonormal basis $\{B_j | j = 1, \dots, k\}$, $k \leq n$, of the observed $\{v_i\}$. Project each of the vectors v_i onto this orthonormal basis using $v_i \approx \sum_{j=1}^k c_{i,j} B_j$, where $c_{i,j} = \langle\langle v_i, B_j \rangle\rangle$. Now, each original shape can simply be represented using the coefficient vector $c_i = \{c_{i,j}\}$. Then, the sample covariance matrix can be computed in the coefficient space as $K = (1/(n-1)) \sum_{i=1}^n c_i c_i^T \in \mathbb{R}^{k \times k}$. One can use the singular value decomposition of K to determine the principal directions of variation in the given data. For example, if $u \in \mathbb{R}^k$ corresponds to a principal singular vector of K , then the corresponding tangent vector in $T_{\bar{f}}(\mathcal{F})$ is given by $v = \sum_{j=1}^k u_j B_j$. One can then map this vector to a surface f using the exponential map, which is used to map points from the tangent space to the representation space. The exponential map must be computed under one of the nonstandard metrics introduced earlier, which is not a simple task. This can be accomplished using a tool called parallel transport, which was derived for the SRNF representation of surfaces by Xie et al. [50]. For brevity, we do not provide details here but rather refer the interested reader to that paper. We also note that when computing the following results we approximated the exponential map using a linear mapping.

Given the mean and covariance, we can impose a Gaussian distribution in the tangent space at the mean shape. This will provide a model, which can be used to generate random shapes. A random tangent vector $v \in T_{\bar{f}}(\mathcal{F})$ can be sampled from the Gaussian distribution using $v = \sum_{j=1}^k z_j \sqrt{S_{jj}} u_j B_j$, where $z_j \stackrel{iid}{\sim} N(0, 1)$, S_{jj} is the variance of the j th principal component, u_j is the corresponding principal singular vector, and B_j is a basis element as defined previously. One can then map this element of the tangent space to a surface shape using the exponential map to obtain a random shape from the Gaussian distribution.

We present an example of computing shape statistics for a toy data set in Figs. 12.6 and 12.7. The data in this example was simulated in such a way that one of the peaks on each surface was already matched perfectly while the position of the second peak was slightly perturbed. All surfaces in this data set are displayed in the top panel of Fig. 12.6. We computed three averages in this example: simple average in the pre-shape space (bottom left of Fig. 12.6), Karcher mean under the SRF pullback metric $\langle\langle \cdot, \cdot \rangle\rangle$ (bottom center of Fig. 12.6), and Karcher mean under the partial elastic metric $\langle\langle\langle \cdot, \cdot \rangle\rangle\rangle$ (bottom right of Fig. 12.6). The pre-shape mean is not a very good summary of the given data. One of the peaks is sharp while the other is averaged out due to misalignment. The Karcher mean under the SRF pullback metric is much better, although it also shows slight averaging out of the second peak. The best representative shape is given by the Karcher mean computed under the partial elastic metric where the two peaks are sharp as in the original data. In Fig. 12.7, we display the two principal directions of variation in the given data computed by those three methods. The result computed in the pre-shape space does not reflect the true variability in the given data. In fact, as one goes in the positive second principal direction, the resulting shapes have three peaks. This result is again

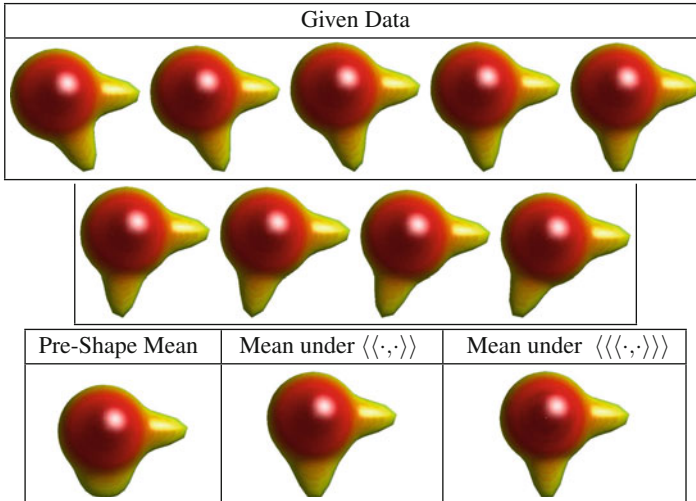


Fig. 12.6 *Top:* given data. *Bottom:* shape means computed (1) in the pre-shape space, (2) under the SRF pullback metric, and (3) under the partial elastic metric

improved under the SRF pullback metric. But, there is still some misalignment, which can be seen in the second principal direction where a wide peak evolves into a thin peak. The best result is observed in the case of the partial elastic metric. Here, all of the variability is contained in the first principal direction of variation where the peak naturally moves without any distortion. Based on the given data, this is the most intuitive summary of variability.

In Fig. 12.8, we show two random samples from the Gaussian distribution defined (1) in the pre-shape space, (2) in the shape space under the SRF pullback metric, and (3) in the shape space under the partial elastic metric. The two random samples in the pre-shape space do not resemble the given data as they both have three peaks. While method (2) produced random samples with two peaks, one can see that in both cases one of the peaks has a strange shape (either too thin or too wide). Method (3) produces the best random samples, which clearly resemble the structure in the given data.

12.4 Classification of Attention Deficit Hyperactivity Disorder

In this last section, we present an application of the described methods to medical imaging: the diagnosis of attention deficit hyperactivity disorder (ADHD) using structural MRI. The results presented here have been previously published in [31, 32, 47]. The surfaces of six left and right brain structures were extracted from

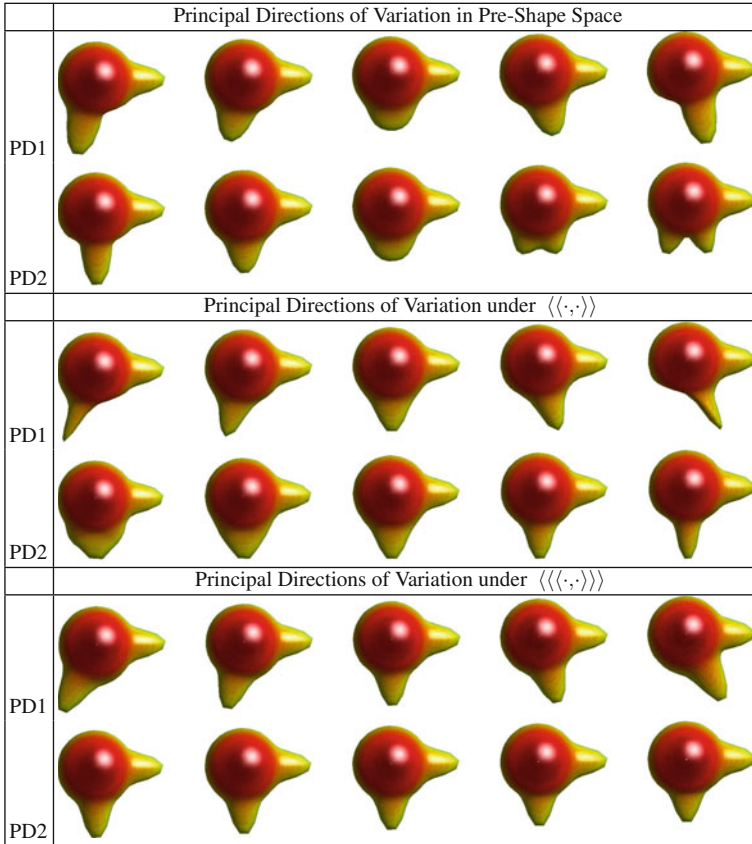


Fig. 12.7 The two principal directions of variation (from -2σ to $+2\sigma$) for (1) pre-shape mean and covariance without optimization over Γ , (2) pullback metric under SRF representation, and (3) partial elastic metric

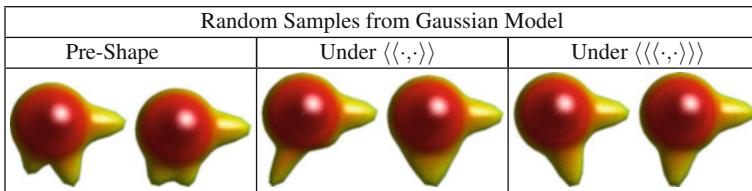


Fig. 12.8 Random samples from Gaussian model computed (1) in the pre-shape space, (2) under the pullback metric under SRF representation, and (3) under the partial elastic metric

T1-weighted brain magnetic resonance images of young adults aged between 18 and 21. These subjects were recruited from the Detroit Fetal Alcohol and Drug Exposure Cohort [7, 23, 24]. Among the 34 subjects studied, 19 were diagnosed

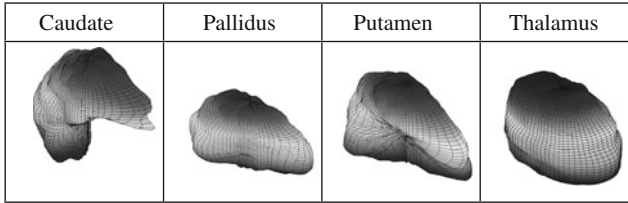


Fig. 12.9 Left subcortical structures used for ADHD classification. (Courtesy of Kurtek et al. [32])

Table 12.1 Classification accuracy (%) for five different techniques

| Method structure (%) | SRNF Gauss [47] | SRF Gauss [31] | SRF NN [32] | Harmonic | ICP [2] | SPHARM-PDM [44] |
|-------------------------|--------------------|-------------------|----------------|----------|------------|--------------------|
| L. Caudate | 67.7 | – | 41.2 | 64.7 | 32.4 | 61.8 |
| L. Pallidus | 85.3 | 82.4 | 76.5 | 79.4 | 67.7 | 44.1 |
| L. Putamen | 94.1 | 88.2 | 82.4 | 70.6 | 61.8 | 50.0 |
| L. Thalamus | 67.7 | – | 58.8 | 67.7 | 35.5 | 52.9 |
| R. Caudate | 55.9 | – | 50.0 | 44.1 | 50.0 | 70.6 |
| R. Pallidus | 76.5 | 67.6 | 61.8 | 67.7 | 55.9 | 52.9 |
| R. Putamen | 67.7 | 82.4 | 67.7 | 55.9 | 47.2 | 55.9 |
| R. Thalamus | 67.7 | – | 58.8 | 52.9 | 64.7 | 64.7 |

with ADHD and the remaining 15 were controls (non-ADHD). Some examples of the left structures used for classification are displayed in Fig. 12.9. In order to distinguish between ADHD and controls, we utilize several methods for comparison: SRNF Gaussian, SRF Gaussian, SRF distance nearest neighbor (NN), iterative closest point (ICP) algorithm, \mathbb{L}^2 distance between fixed surface parameterizations (Harmonic), and SPHARM-PDM. The classification performance for all methods is reported in Table 12.1. The results suggest that the Riemannian approaches based on parameterization-invariant metrics perform best in this setting due to improved matching across surfaces. Furthermore, the use of Gaussian probability models outperforms standard distance-based classification.

12.5 Summary

This chapter describes two Riemannian frameworks for statistical shape analysis of 3D objects. The important feature of these methods is that they define reparameterization-invariant Riemannian metrics on the space of parameterized surfaces. The first framework develops the metric by pulling back the \mathbb{L}^2 metric from the space of square-root function representations of surfaces. The main drawbacks of this method are the lack of interpretability of the resulting Riemannian metric and the lack of invariance to translations. The second framework starts with a full

elastic metric on the space of parameterized surfaces. This metric is then restricted, resulting in a partial elastic metric, which becomes the simple \mathbb{L}^2 metric under the square-root normal field representation of surfaces. This metric has a nice interpretation in terms of the amount of stretching and bending that is needed to deform one surface into another. We show examples of geodesic paths and distances between complex surfaces for both cases. Given the ability to compute geodesics in the shape space, we define the first two moments, the Karcher mean and covariance, and use them for random sampling from Gaussian-type shape models. Finally, we showcase the applicability of these methods in an ADHD medical study.

Acknowledgements This work was partially supported by a grant from the Simons Foundation (#317865 to Eric Klassen).

References

1. Abe K, Erbacher J (1975) Isometric immersions with the same Gauss map. *Math Ann* 215(3):197–201
2. Almhdie A, Léger C, Deriche M, Lédée R (2007) 3D registration using a new implementation of the ICP algorithm based on a comprehensive lookup matrix: application to medical imaging. *Pattern Recogn Lett* 28(12):1523–1533
3. Bauer M, Bruveris M (2011) A new Riemannian setting for surface registration. In: International workshop on mathematical foundations of computational anatomy, pp 182–193
4. Beg M, Miller M, Trounev A, Younes L (2005) Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int J Comput Vis* 61(2):139–157
5. Bouix S, Pruessner JC, Collins DL, Siddiqi K (2001) Hippocampal shape analysis using medial surfaces. *NeuroImage* 25:1077–1089
6. Brechbühler C, Gerig G, Kübler O (1995) Parameterization of closed surfaces for 3D shape description. *Comput Vis Image Underst* 61(2):154–170
7. Burden MJ, Jacobson JL, Westerlund AJ, Lundahl LH, Klorman R, Nelson CA, Avison MJ, Jacobson SW (2010) An event-related potential study of response inhibition in ADHD with and without prenatal alcohol exposure. *Alcohol Clin Exp Res* 34(4):617–627
8. Cates J, Meyer M, Fletcher P, Whitaker R (2006) Entropy-based particle systems for shape correspondence. In: International workshop on mathematical foundations of computational anatomy, pp 90–99
9. Christensen GE, Johnson H (2001) Consistent image registration. *IEEE Trans Med Imaging* 20(7):568–582
10. Clarke B (2010) The metric geometry of the manifold of Riemannian metrics over a closed manifold. *Calc Var Partial Differ Equ* 39(3):533–545
11. Csernansky JG, Wang L, Jones DJ, Rastogi-Cru D, Posener G, Heydebrand JA, Miller JP, Grenander U, Miller MI (2002) Hippocampal deformities in schizophrenia characterized by high dimensional brain mapping. *Am J Psychiatry* 159(12):1–7
12. Davatzikos C, Vaillant M, Resnick S, Prince JL, Letovsky S, Bryan RN (1996) A computerized method for morphological analysis of the corpus callosum. *J Comput Assist Tomogr* 20:88–97
13. Davies RH, Twining CJ, Cootes TF, Taylor CJ (2010) Building 3-D statistical shape models by direct optimization. *IEEE Trans Med Imaging* 29(4):961–981
14. DeWitt BS (1967) Quantum theory of gravity. I. The canonical theory. *Phys Rev* 160(5):1113–1148
15. Dryden IL, Mardia KV (1998) Statistical shape analysis. Wiley, New York

16. Ebin D (1970) The manifold of Riemannian metrics. In: *Symposia in pure mathematics* (AMS), vol 15, pp 11–40
17. Gerig G, Styner M, Shenton ME, Lieberman JA (2001) Shape versus size: improved understanding of the morphology of brain structures. In: *MICCAI*, pp 24–32 (2001)
18. Gil-Medrano O, Michor PW (1991) The Riemannian manifold of all Riemannian metrics. *Q J Math* 42:183–202
19. Gorczowski K, Styner M, Jeong JY, Marron JS, Piven J, Hazlett HC, Pizer SM, Gerig G (2010) Multi-object analysis of volume, pose, and shape using statistical discrimination. *IEEE Trans Pattern Anal Mach Intell* 32(4):652–666
20. Grenander U, Miller MI (1998) Computational anatomy: an emerging discipline. *Q Appl Math LVI*(4):617–694
21. Gu X, Vemuri BC (2004) Matching 3D shapes using 2D conformal representations. In: *MICCAI*, pp 771–780
22. Gu X, Wang S, Kim J, Zeng Y, Wang Y, Qin H, Samaras D (2007) Ricci flow for 3D shape analysis. In: *IEEE international conference on computer vision*, pp 1–8
23. Jacobson SW, Jacobson JL, Sokol RJ, Chiodo LM, Corobana R (2004) 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
24. Jacobson SW, Jacobson JL, Sokol RJ, Martier SS, Chiodo LM (1996) New evidence for neurobehavioral effects of in utero cocaine exposure. *J Pediatr* 129:581–590
25. Jermyn IH, Kurtek S, Klassen E, Srivastava A (2012) Elastic shape matching of parameterized surfaces using square root normal fields. In: *European conference on computer vision*, pp 804–817
26. Joshi S, Miller M, Grenander U (1997) On the geometry and shape of brain sub-manifolds. *Pattern Recogn Artif Intell* 11:1317–1343
27. Joshi SH, Klassen E, Srivastava A, Jermyn IH (2007) A novel representation for Riemannian analysis of elastic curves in \mathbb{R}^n . In: *IEEE conference on computer vision and pattern recognition*, pp 1–7
28. van Kaick O, Zhang H, Hamarneh G, Cohen-Or D (2010) A survey on shape correspondence. In: *Eurographics state-of-the-art report*, pp 1–24
29. Kelemen A, Szekely G, Gerig G (1999) Elastic model-based segmentation of 3D neurological data sets. *IEEE Trans Med Imaging* 18(10):828–839
30. Kilian M, Mitra NJ, Pottmann H (2007) Geometric modeling in shape space. *ACM Trans Graph* 26(3):64
31. Kurtek S, Klassen E, Ding Z, Avison MJ, Srivastava A (2011) Parameterization-invariant shape statistics and probabilistic classification of anatomical surfaces. In: *Information processing in medical imaging*, pp 147–158
32. Kurtek S, Klassen E, Ding Z, Jacobson SW, Jacobson JL, Avison MJ, Srivastava A (2011) Parameterization-invariant shape comparisons of anatomical surfaces. *IEEE Trans Med Imaging* 30(3):849–858
33. Kurtek S, Klassen E, Ding Z, Srivastava A (2010) A novel Riemannian framework for shape analysis of 3D objects. In: *IEEE conference on computer vision and pattern recognition*, pp 1625–1632
34. Kurtek S, Klassen E, Gore J, Ding Z, Srivastava A (2012) Elastic geodesic paths in shape space of parameterized surfaces. *IEEE Trans Pattern Anal Mach Intell* 34(9):1717–1730
35. Kurtek S, Klassen E, Gore JC, Ding Z, Srivastava A (2011) Classification of mathematics deficiency using shape and scale analysis of 3D brain structures. In: *SPIE medical imaging: image processing*, vol 7962
36. Kurtek S, Samir C, Ouchchane L (2014) Statistical shape model for simulation of realistic endometrial tissue. In: *International conference on pattern recognition applications and methods*
37. Kurtek S, Srivastava A, Klassen E, Ding Z (2012) Statistical modeling of curves using shapes and related features. *J Am Stat Assoc* 107(499):1152–1165

38. Kurtek S, Srivastava A, Klassen E, Laga H (2013) Landmark-guided elastic shape analysis of spherically-parameterized surfaces. *Comput Graph Forum* 32(2):429–438
39. Malladi R, Sethian JA, Vemuri BC (1996) A fast level set based algorithm for topology-independent shape modeling. *J Math Imaging Vis* 6:269–290
40. Mio W, Srivastava A, Joshi SH (2007) On shape of plane elastic curves. *Int J Comput Vis* 73(3):307–324
41. Samir C, Kurtek S, Srivastava A, Canis M (2014) Elastic shape analysis of cylindrical surfaces for 3D/2D registration in endometrial tissue characterization. *IEEE Trans Med Imaging* 33(5):1035–1043
42. Srivastava A, Klassen E, Joshi SH, Jermyn IH (2011) Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans Pattern Anal Mach Intell* 33(7):1415–1428
43. Srivastava A, Turaga P, Kurtek S (2012) On advances in differential-geometric approaches for 2D and 3D shape analyses and activity recognition. *Image Vis Comput* 30(6–7):398–416
44. Styner M, Oguz I, Xu S, Brechbühler C, Pantazis D, Levitt J, Shenton ME, Gerig G (2006) Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *Insight J* 1071:242–250
45. Tagare H, Groisser D, Skrinjar O (2009) Symmetric non-rigid registration: a geometric theory and some numerical techniques. *J Math Imaging Vis* 34(1):61–88
46. Vaillant M, Glaunes J (2005) Surface matching via currents. In: *Information processing in medical imaging*, pp 381–392
47. Xie Q, Jermyn IH, Kurtek S, Srivastava A (2014) Numerical inversion of SRNFs for efficient elastic shape analysis of star-shaped objects. In: *European conference on computer vision*, pp 485–499
48. Xie Q, Kurtek S, Christensen GE, Ding Z, Klassen E, Srivastava A (2012) A novel framework for metric-based image registration. In: *Workshop on biomedical image registration*, pp 276–285
49. Xie Q, Kurtek S, Klassen E, Christensen GE, Srivastava A (2014) Metric-based pairwise and multiple image registration. In: *European conference on computer vision*, pp 236–250
50. Xie Q, Kurtek S, Le H, Srivastava A (2013) Parallel transport of deformations in shape space of elastic surfaces. In: *IEEE international conference on computer vision*, pp. 865–872
51. Younes L (1998) Computable elastic distance between shapes. *SIAM J Appl Math* 58(2):565–586