

A Novel Riemannian Framework for Shape Analysis of 3D Objects

Sebastian Kurtek
Department of Statistics
Florida State University, Tallahassee, FL
skurtek@stat.fsu.edu

Zhaohua Ding
Institute of Imaging Science
Vanderbilt University, Nashville, TN
zhaohua.ding@vanderbilt.edu

Eric Klassen
Department of Mathematics
Florida State University, Tallahassee, FL
klassen@math.fsu.edu

Anuj Srivastava
Department of Statistics
Florida State University, Tallahassee, FL
anuj@stat.fsu.edu

Abstract

In this paper we introduce a novel Riemannian framework for shape analysis of parameterized surfaces. We derive a distance function between any two surfaces that is invariant to rigid motion, global scaling, and re-parametrization. It is the last part that presents the main difficulty. Our solution to this problem is twofold: (1) we define a special representation, called a q -map, to represent each surface, and (2) we develop a gradient-based algorithm to optimize over different re-parameterizations of a surface. The second step is akin to deforming the mesh on a fixed surface to optimize its placement. (This is different from the current methods that treat the given meshes as fixed.) Under the chosen representation, with the \mathbb{L}^2 metric, the action of the re-parametrization group is by isometries. This results in, to our knowledge, the first Riemannian distance between parameterized surfaces to have all the desired invariances. We demonstrate this framework with several examples using some toy shapes, and real data with anatomical structures, and cropped facial surfaces. We also successfully demonstrate clustering and classification of these objects under the proposed metric.

1. Introduction

Shape analysis can play a very important role in image analysis and computer vision, with applications in medical diagnostics, bioinformatics, robotics, target recognition, and many more branches of science. Figure 1 shows a number of anatomical objects in a human brain that are represented by meshes on their surfaces and our goal is to analyze their shapes. There has been an increasing interest




Figure 1. A display of anatomical structures in the human brain whose shapes are studied by treating them as continuous surfaces.

in using a Riemannian framework for shape analysis of objects. One reason for its increasing popularity is the breadth of the tools that it offers. On one hand it allows us to remove all shape-preserving transformations from the representation space, using the notion of quotient spaces and, on the other, it provides ideas and algorithms for computing statistics (sample means, covariances, etc.) of shapes. However, most of this work is limited to shape analysis of curves (see [11, 7] and references therein.) There are very few papers that study shapes of continuous 2D surfaces in a similar fashion. Just to clarify this further, there are many papers in the literature that study 3D objects in other ways e.g. by choosing landmarks on them and performing Kendall-type analysis [3], applying thin plate splines, or ICP [1]; using medial surfaces [2]; using level sets [9]; deforming the volumes (background spaces) containing these objects [5]. However, very few of them study the shapes of parameterized surfaces directly, in a Riemannian way that is analogous to shape analysis of parameterized curves. A recent paper by Fuchs et al. [4] considers both the curves and




Figure 2. Same right pallidum surface with four different parameterizations. We seek a shape metric that will result in zero distance between them.

their interiors, but falls short of studying 3D objects. One exception is Kilian et al. [8] which represents parameterized surfaces by discrete triangulated meshes, assumes a Riemannian metric on the space of such meshes, and computes geodesic paths between given meshes. While this paper represents a major progress, it has a limitation in that it *assumes* the correspondence between points across meshes. That is, *we need to know beforehand which point on one mesh matches with which point on the second mesh*. A more formal way of saying this is that this method works with fixed parameterizations of surfaces, rather than analyzing them. In contrast, we would like to remove the re-parametrization variability so that different surfaces with the same shape but different parameterizations (e.g. the four surfaces shown in Figure 2) will have zero distance between them. We can obtain the data for parameterized surfaces in several different ways. For example, for cropped facial surfaces, we obtain the mesh directly from the laser scanner. For anatomical volume data, we utilize the marching cube algorithm or the moving least-squares technique to reconstruct the surface.

How is this problem (of removing parametrization) handled in the shape analysis of curves? Several papers solve it in a similar way. They define a new representation of curves such that under the standard \mathbb{L}^2 metric the re-parametrization of curves preserves distances. Let $\beta_1, \beta_2 : [0, 1] \rightarrow \mathbb{R}^n$ be two parameterized curves, and let Γ be the re-parametrization group (the set of all corner preserving diffeomorphisms of $[0, 1]$). Then, it is easy to see that $\|\beta_1 - \beta_2\| \neq \|\beta_1 \circ \gamma - \beta_2 \circ \gamma\|$ for $\gamma \in \Gamma$ in general. Joshi et al [7] define a function $q(t) = \dot{\beta}(t) / \sqrt{\|\dot{\beta}(t)\|}$ (for $n = 2$ there are several formulae, but, so far, this is the only one for $n > 2$) that is used to represent a curve β , with the following properties. First, note that the q -function for a re-parameterized curve $\beta \circ \gamma$ is given by $(q, \gamma) = (q \circ \gamma) \sqrt{\dot{\gamma}}$. Then, note that $\|q_1 - q_2\| = \|(q_1, \gamma) - (q_2, \gamma)\|$ for all

$\gamma \in \Gamma$. This way, one can remove the re-parametrization group and define a distance between shapes of curves using: $\inf_{\gamma \in \Gamma} \|q_1 - (q_2, \gamma)\|$. We emphasize that this is a proper distance in the shape space of curves, i.e. it satisfies all three properties of a distance, including the triangle inequality.

Our goal is to develop a similar representation of surfaces that allows invariant shape analysis. Seemingly simple, this problem is complicated by the following issue. Unlike curves there is no ordering of points on surfaces. One consequence is that the space of re-parameterizations of surfaces is much larger (i.e. the group of re-parameterizations of $[0, 1]^2$ versus that for $[0, 1]$). Secondly, the choice of the q -map is not obvious. In this paper we propose a novel q -map for a surface such that a re-parametrization of two surfaces preserves the \mathbb{L}^2 distance between their q -maps. This leads to the definition of a quotient space of surfaces, modulo re-parametrization group, and the \mathbb{L}^2 metric makes this quotient space a Riemannian manifold. The geodesic distances on this quotient space are used to compare, match, and align shapes of any two surfaces. The matching part requires solving an optimization problem on the group of all re-parameterizations, an infinite-dimensional space. We present a novel gradient technique for finding solutions, albeit local, to this optimization.

The rest of this paper is organized as follows. Section 2 introduces the q -representation of surfaces for analyzing their shapes, while Section 3 lays out a gradient-based technique for matching surfaces, by optimizing over the re-parametrization group. Section 4 illustrates these ideas using some examples from computational anatomy and facial biometrics.

2. Mathematical Framework

Let S denote a 2D smooth surface with genus zero. Several past papers have represented closed surfaces as real-valued functions on \mathbb{S}^2 (for instance, using spherical harmonics [10]), or as embeddings f of \mathbb{S}^2 in \mathbb{R}^3 , and have applied the standard \mathbb{L}^2 metric for comparing shapes. We investigate that idea more carefully here. Let D be a domain of parametrization; we will consider $D = [0, 1]^2$, \mathbb{S}^2 and the unit disk as possible parametrization domains. For example, we will use $D = [0, 1]^2$ for quadrilateral surfaces, $D = \mathbb{S}^2$ for closed surfaces and $D = \text{unit disk}$ for hemispherical surfaces. Then, represent the surface S with its embedding $f : D \rightarrow \mathbb{R}^3$; f is also called a *parametrization* of S . Let the set of all such parameterized surfaces be $\mathcal{F} = \{f : D \mapsto \mathbb{R}^3 \mid \int_D \|f(s)\|^2 ds < \infty\}$, where ds is the standard Lebesgue measure on D . For any function $f \in \mathcal{F}$, the tangent space $T_f(\mathcal{F})$ is \mathcal{F} itself. We choose the natural Riemannian structure on \mathcal{F} : for any two elements $v_1, v_2 \in T_f(\mathcal{F})$, define an inner product: $\langle v_1, v_2 \rangle = \int_D \langle v_1(s), v_2(s) \rangle ds$. (The integrand is simply the Euclidean inner product in \mathbb{R}^3 .) The result-

ing \mathbb{L}^2 distance between any two points $f_1, f_2 \in \mathcal{F}$ is $(\int_D \|f_1(s) - f_2(s)\|^2 ds)^{1/2}$, and the geodesic path connecting them in \mathcal{F} is a “straight line”: $\psi(t) = tf_2 + (1 - t)f_1$. One can represent surfaces as elements of \mathcal{F} and use the \mathbb{L}^2 distance to compare shapes of surfaces. What is the problem with this idea? The problem is that this metric is not invariant to re-parameterizations of surfaces!

Let Γ be the set of all diffeomorphisms of D ; this set will act as the re-parameterization group for surfaces. The natural action of Γ on \mathcal{F} is on the right by composition: for a $\gamma \in \Gamma$, $f \in \mathcal{F}$, the re-parameterized surface is given by $f \circ \gamma$. We check the isometry condition:

$$\begin{aligned} d(f_1 \circ \gamma, f_2 \circ \gamma) &= \left(\int_D \|f_1(\gamma(s)) - f_2(\gamma(s))\|^2 ds \right)^{1/2} \\ &= \left(\int_D \|f_1(\tilde{s}) - f_2(\tilde{s})\|^2 J_\gamma(\tilde{s})^{-1} d\tilde{s} \right)^{1/2} \end{aligned}$$

which is not equal to $d(f_1, f_2)$. Here $J_\gamma : D \rightarrow \mathbb{R}$ is the determinant of the Jacobian matrix of γ at a point $s \in D$, when this matrix is written with respect to an orthonormal basis of $T_s(D)$. The inequality comes from the fact that γ , in general, is not area preserving and hence J_γ is not one at all points. This lack of isometry makes the natural action of Γ on \mathcal{F} difficult to use in shape analysis of surfaces. This means that the distance between any two surfaces will not be the same if they are re-parameterized by the same element of Γ . One solution is to restrict to only those re-parameterizations that are area-preserving or some subset of these [6]. However, this is quite restrictive and may not be able to provide a good matching of surfaces. We suggest an alternate approach. Define a new mathematical representation of a surface as follows.

Definition 1. Given a differentiable surface f , define $\kappa : D \rightarrow \mathbb{R}_+$ to be the “volume multiplication factor” of f at s . To be more precise, let $\kappa(s) = \|df_s(u) \times df_s(v)\|$, where $\{u, v\}$ is an orthonormal basis of $T_s(D)$. Now, define $q : D \rightarrow \mathbb{R}^3$ by $q(s) = \sqrt{\kappa(s)}f(s)$.

The factor $\kappa(s)$ is the ratio of infinitesimal areas of the surface at s and the domain at s . For a point $s \in D$, the differential $df_s : T_s(D) \rightarrow \mathbb{R}^3$ is a linear mapping between vector spaces.

In our approach each surface S will be represented by its q -map. The set of all q -maps is also taken to be \mathcal{F} . To make it a Riemannian manifold, we will use the \mathbb{L}^2 metric. As noted earlier, geodesics in \mathcal{F} under the \mathbb{L}^2 metric are simply straight lines. In order for the shape analysis to be invariant to certain group actions, we normalize the surfaces as follows. The translation of surfaces is easily taken care off by normalizing their locations: $f_{centered}(s) = f(s) - \frac{\int_D f(s)\kappa(s)ds}{\int_D \kappa(s)ds}$. The total area of a

surface is given by $\int_D \kappa(s)ds$. We use this fact to manually scale all of the surfaces by fixing the surface area to 1: $f_{scaled}(s) = \frac{f(s)}{\sqrt{\int_D \kappa(s)ds}}$. The remaining groups – rotation and re-parameterization – are dealt with differently, by removing them algebraically from the representation.

1. **Rotation Group, $SO(3)$:** The rotation group $SO(3)$ acts on \mathcal{F} , $SO(3) \times \mathcal{F} \rightarrow \mathcal{F}$ according to $(O, q) = Oq$, for $O \in SO(3)$ and $q \in \mathcal{F}$. It is easy to check that the action of $SO(3)$ preserves the \mathbb{L}^2 distance on \mathcal{F} , i.e. $\|Oq_1 - Oq_2\| = \|q_1 - q_2\|$ for all $O \in SO(3)$, $q_1, q_2 \in \mathcal{F}$.
2. **Re-Parameterization Group, Γ :** Given a $\gamma \in \Gamma$ and a parameterized surface $f \in \mathcal{F}$, what is the q -map of the re-parameterized surface? This is given by the following group action.

Definition 2. Define a right action $\mathcal{F} \times \Gamma \rightarrow \mathcal{F}$ by $(q, \gamma) = \sqrt{J_\gamma}(q \circ \gamma)$.

An important fact about the map $f \mapsto q$ is that if we re-parameterize a surface by γ and then obtain its q representation (Definition 1), or if we obtain its q representation (Definition 1) and then act by γ (Definition 2), the result will be the same.

Proposition 1. The re-parameterization group Γ acts on \mathcal{F} by isometries. That is, for any $q_1, q_2 \in \mathcal{F}$ and $\gamma \in \Gamma$, we have $\|q_1 - q_2\| = \|(q_1, \gamma) - (q_2, \gamma)\|$.

Proposition 2. The actions of Γ and $SO(3)$ on \mathcal{F} commute.

Since, the actions of $SO(3)$ and Γ commute we are able to define an action of the product of the groups on \mathcal{F} . We are also able to define the quotient space $\mathcal{S} \equiv \mathcal{F}/(\Gamma \times SO(3))$ as the shape space. The orbits of a surface are given by:

$$[q] = \{(Oq, \gamma) | O \in SO(3), \gamma \in \Gamma\} \quad (1)$$

and the set of all $[q]$ is \mathcal{S} . Furthermore, since the product group of $SO(3)$ and Γ acts on \mathcal{C} by isometries, geodesics in the pre-shape space \mathcal{C} , which are perpendicular to the orbits of q_1 and q_2 map to geodesics in the quotient space \mathcal{S} . We can define the distance between orbits as follows,

Definition 3. Define a Riemannian distance function between any two surfaces S_1 and S_2 , represented by their q -maps q_1 and q_2 , as

$$d([q_1], [q_2]) \equiv \inf_{O \in SO(3), \gamma \in \Gamma} \|q_1 - (Oq_2, \gamma)\|. \quad (2)$$

Remark: This optimization over Γ is similar to the deformable template approach in image analysis where one solves a problem of the type

$$P(I_1, I_2) = \min_{\gamma \in \Gamma} \|I_1 - I_2 \circ \gamma\|^2 + \text{penalty on } \gamma.$$

However, $P(I_1, I_2)$ is not a distance between images; it is not even symmetric.

Note that in Eqn. 2 we only need to apply γ , O to q_2 , not also to q_1 . This distance will serve as a measure of differences between shapes of two surfaces $d(f_1, f_2) \equiv d([q_1], [q_2])$. The computation of this distance requires solving the joint optimization problem on $SO(3) \times \Gamma$, of which the optimization on Γ is more difficult. In fact, for a fixed $\gamma \in \Gamma$, the minimization over $SO(3)$ can be performed directly using the Procrustes analysis. Let \tilde{q}_2 denote (q_2, γ) ; then, the optimal value of O is obtained as follows. Compute the 3×3 matrix $A = \int_D q_1(s) \tilde{q}_2(s)^T ds$. Using the singular value decomposition $A = U \Sigma V^T$, we can define the optimal rotation as $O^* = VU^T$ (if the determinant of A is negative, the last column of V changes sign). If we can address optimization over Γ , then our approach will be able to alternate between the minimizations over $SO(3)$ and Γ and converge to a solution.

3. Gradient Over Re-Parametrization Group

We consider the optimization problem in Eqn. 2 over Γ and use a gradient approach to solve it. It is this optimization that makes this approach novel. Although this provides only a local solution, we do not have any alternative at this stage. Of course, the initialization of the algorithm becomes important and we will discuss some strategies later. We remark that the dynamic programming algorithm that is important in matching of curves is not directly applicable to the surface matching problem.




Figure 3. Pictorial depiction of the gradient descent iteration.

First we explain this iterative process using Figure 3, and then present the mathematical details. At the current iteration, let the current re-parametrization of the second surface result in $r \in [q_2]$. Let $T_r([q_2])$ be the space of tangents to the Γ -orbit $[q_2]$ at r and set $v = \frac{q_1 - r}{\|q_1 - r\|}$ in $T_r([q_2])$. Actually v is the gradient of the cost function $E \equiv \|q_1 - r\|$, given in Eqn. 2, with respect to r . If the projection of v along the orbit $[q_2]$ is zero, then the straight line from q_1 to r is perpendicular to the orbit $[q_2]$, and indeed all other orbits; this is the geodesic distance between the elements of $[q_1]$ and $[q_2]$. Otherwise, we update r in the direction of this projection. An integral step here is to compute the projection of v on $T_r([q_2])$. This step is performed numerically: we deter-

mine an orthonormal basis of a finite-dimensional subspace of $T_r([q_2])$ and use those basis elements to approximate the projection of v . These basis elements, in turn, are first derived in the space $T_{\gamma_{id}}(\Gamma)$ and then transferred to $T_r([q_2])$ using an appropriate mapping.

1. **Tangent Space $T_{\gamma_{id}}(\Gamma)$:** Γ is a Lie group with the composition as the group operation and the identity mapping $\gamma_{id}(s) = s$ as the identity element. The tangent space of Γ at γ_{id} is a set of smooth vector fields on D that are tangential to D at its boundaries. We need an orthonormal basis for $T_{\gamma_{id}}(\Gamma)$. These basis elements are defined so that they preserve the boundaries of the domain D . Although such boundary restrictions are not present for $D = \mathbb{S}^2$, this case is complicated by itself. We define the basis vector fields for \mathbb{S}^2 using spherical harmonics and its gradients as follows. Take the gradients of these harmonics to obtain a set of vector fields on \mathbb{S}^2 . Add to this set another set of vector fields, each of which is obtained by rotating a vector field in the original set by 90° , and we get the full \mathbb{L}^2 basis of $T_{\gamma_{id}}(\Gamma)$ for $D = \mathbb{S}^2$. We do not present the construction of these vector fields in detail, but we show some examples in Figure 4 for each of the three domains. We will use $\{b_i : D \rightarrow \mathbb{R}^2 | i = 1, 2, \dots\}$ to denote the resulting orthonormal basis of $T_{\gamma_{id}}(\Gamma)$.
2. **Differential of Group Action:** Continue assuming that r is the current estimate in $[q_2]$ and define a map $\phi : \Gamma \rightarrow [q_2]$ given by $\phi(\gamma) = (r, \gamma)$. We need to determine the derivative of this function at the identity, $d\phi : T_{\gamma_{id}}(\Gamma) \rightarrow T_r([q_2])$. To set this up, let $b \in T_{\gamma_{id}}(\Gamma)$ and separate r into its components $r = (r_1, r_2, r_3)$, where each r_j is a real-valued function on D . Then $\phi_*(b)$ will also have 3 components, given by the formula:

$$\phi_*(b)_j = \frac{1}{2} (\nabla \cdot b) r_j + \nabla r_j \cdot b \quad j = 1, 2, 3. \quad (3)$$

This formula provides a way for mapping the elements of $T_{\gamma_{id}}(\Gamma)$ into $T_r([q_2])$, and thus taking an orthogonal basis of the first space into a basis (not necessarily orthogonal) for the second space.

We can now state the gradient of the cost function. For the cost function $E(\gamma) = \|q_1 - (r, \gamma)\|$, its gradient with respect to γ at γ_{id} is given by:

$$d\gamma = \sum_{i=1}^{\infty} \langle v, \phi_*(b_i) \rangle b_i \in T_{\gamma_{id}}(\Gamma). \quad (4)$$

This gradient provides the incremental update of γ and sets up the following gradient algorithm.

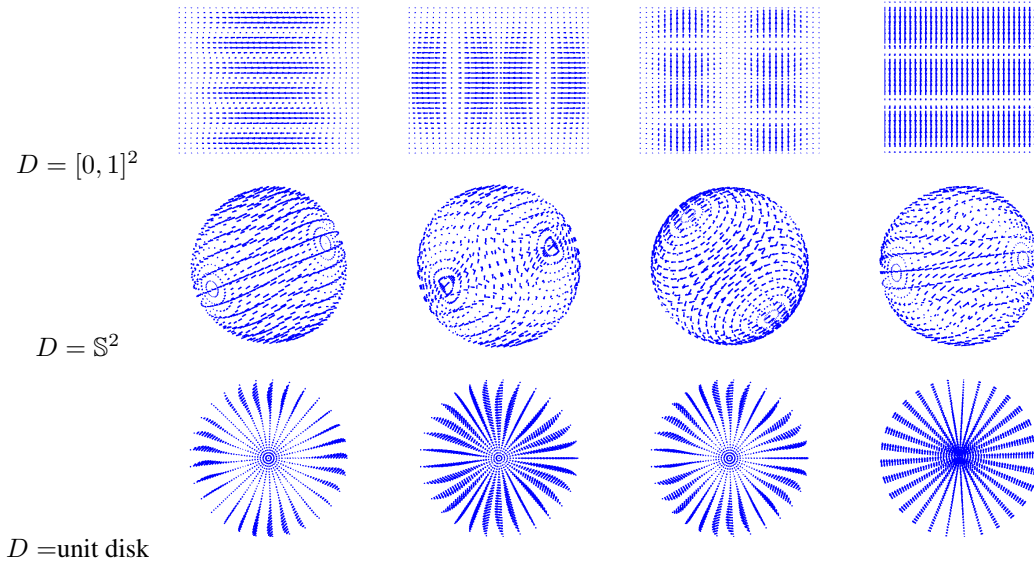


Figure 4. Examples of basis elements of $T_{\gamma_{id}}(\Gamma)$.

Algorithm 1. Let $r \in [q_2]$ be the current estimate of the element nearest to q_1 . Also, let $\{b_i, i = 1, \dots, k\}$ be the orthonormal basis for $T_{\gamma_{id}}(\Gamma)$ for some k . The choice of k is dictated by the computational resources versus the desired accuracy.

1. Compute the tangent vector $v = (q_1 - r)/\|q_1 - r\|$.
2. Compute $\phi_*(b_i)$ for each i using Eqn. 3.
3. Estimate the gradient vector field $d\gamma$ using Eqn. 4.
4. Update the parametrization: $\gamma_{new} = \gamma_{id} + \epsilon d\gamma$, for some step size $\epsilon > 0$.
5. Re-parameterize the second surface f_2 according to $f_2 \circ \gamma_{new}$ and obtain a new r using Defn. 1. Return to step 1. Stop when $\|d\gamma\|$ becomes sufficiently small.

Initialization: An important issue in a gradient search is to initialize the program. We discuss the individual domains separately. For $D = [0, 1]^2$, there are four natural choices for the initial parametrization of the surface f_2 : that surface has four corners and four ways to match the origin $(0, 0) \in D$ to a corner. We try all four possibilities and select the one with the smallest cost. For $D = \mathbb{S}^2$, the problem of initialization is not as simple. Note that Γ in this case contains the compact group $SO(3)$. This set of rigid re-parameterizations corresponds to all different placements of the north pole on \mathbb{S}^3 and all different rotations around that axis. Although $SO(3)$ is much smaller than Γ , it is still an infinite set. Let H denote the largest (irreducible) finite subgroup of $SO(3)$; it is the group of

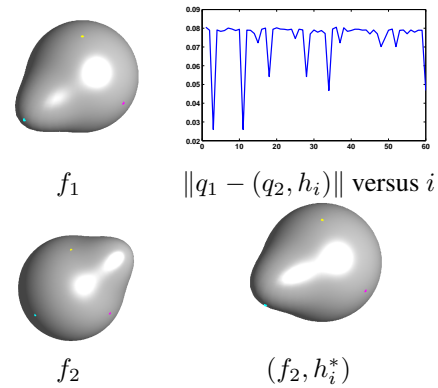


Figure 5. Search over H for initializing Algorithm 1.

symmetries of the dodecahedron, and has 60 elements, each of order 2, 3, or 5. Denote by h_1, h_2, \dots, h_{60} the elements of H ; each of them acts on \mathcal{F} according to Definition 2, i.e. $(q, h_i) = \sqrt{J_{h_i}}(q \circ h_i)$. In case of elements of H , the term J_{h_i} is always one since they are area preserving transformations of \mathbb{S}^2 . So we can search over H by solving for:

$$\hat{i} = \operatorname{argmin}_{i=1,2,\dots,60} \|q_1 - (r, h_i)\|. \quad (5)$$

Shown in Figure 5 is an example of this idea. The top left panel shows the surface f_1 and the bottom left shows f_2 . The cost for 60 different re-parameterizations of f_2 is shown in top right, and the best one is applied to the surface f_2 in bottom right.

4. Experimental Results

In this section, we describe some experimental results to demonstrate the use of this Riemannian distance. Two types of results are presented. First, we show several examples of computing Riemannian distances between the shapes of objects; these objects are toy models, anatomical structures, and cropped facial surfaces. Second, we demonstrate experiments on clustering of 3D anatomical parts and cropped facial surfaces using our shape distance.

4.1. Shape Comparisons of Surfaces

In this section we will present results of surface matching and comparisons using our Riemannian framework. It is important to keep in mind that the distance we provide is a proper Riemannian distance between surfaces. In order to display surface matching we use the following color scheme: we shade the first surface f_1 with a certain color map and then map those colors to the corresponding points of f_2 using the optimal γ^* . Therefore, if the same colors correspond to similar features (e.g. peaks or valleys), it indicates that the matching is good. In addition, in each of these cases we provide a final deformation field on D and the decrease in the cost function E that results from applying Algorithm 1. We display the deformation field as a vector field on the domain D .

In the first example, shown in Figure 6, we compare two parameterized quadrilateral surfaces f_1 and f_2 , when f_2 is in fact a re-parametrization of f_1 , i.e. $f_2 = (f_1 \circ \gamma)$ for some γ . In this case, we expect the resulting Riemannian distance to be zero, or numerically close to zero depending upon the basis size k . As the evolution of the cost function E suggests, the distance $d([q_1], [q_2])$ is found to be 0.0094, almost a 90% decrease in the cost due to gradient iterations. One can see the deformation of the mesh on f_2 in such a way that the new mesh in $(f_2 \circ \gamma^*)$ matches that of f_1 . We check if our distance is symmetric and find that $d([q_2], [q_1]) = 0.0092$.

In order to analyze the tradeoff between computational cost and accuracy the following table summarizes the results obtained using 100 gradient iterations and 900 sample points on meshes (k refers to the number of basis elements used). There is a clear trade-off between the number of basis elements used for optimization over Γ , the accuracy of the distance obtained, and the computational cost. The increase in computational cost is approximately linear with respect to the number of basis elements used.

k	16	48	96	160	240
$d([q_1], [q_2])$	0.0687	0.0496	0.0094	0.0089	0.0088
Comp. Cost (sec)	19.00	28.85	60.44	143.00	311.89

As another toy example, we consider a dumbbell (as f_1) and a unit sphere (as f_2) and the domain D is \mathbb{S}^2 . The

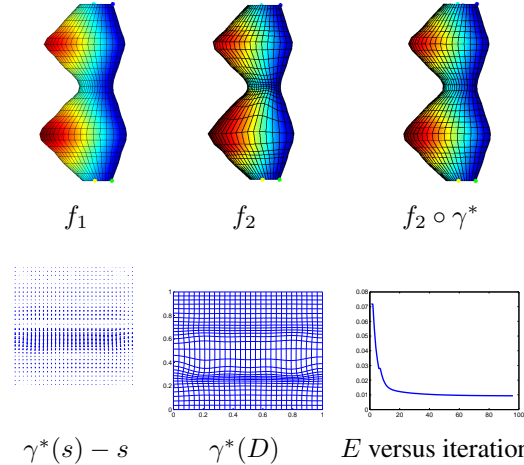


Figure 6. Top: Initial and final matching between surface 1 and a re-parametrization of surface 1. Bottom: Deformation field, final parametrization grid and cost function. $d([q_1], [q_2]) = 0.0094$

change in the shading from f_2 to $(f_2 \circ \gamma^*)$ shows the deformation of the mesh on f_2 to best match the dumbbell shape. We also show the magnitude of the deformation field drawn on f_2 . The light regions denote areas where the grid deformed the most, while the dark regions denote areas of low grid deformation. The decrease in the cost function is also significant - from 0.235 to 0.1748.

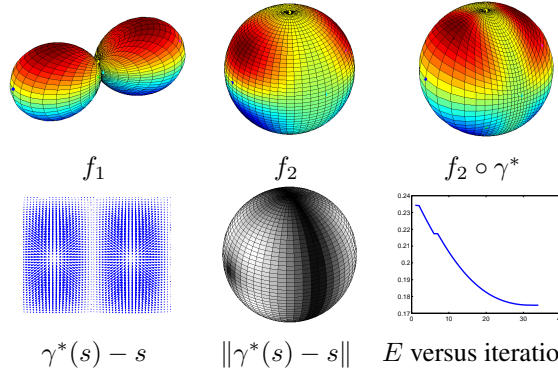


Figure 7. Top: Initial and final matching between a dumbbell and a sphere. Bottom: Deformation and cost function. $d([q_1], [q_2]) = 0.1748$

In the next example, we take a look at some subcortical structures. We compare the shapes of left pallidum surfaces for two subjects. Since these boundaries form closed surfaces, we use $D = \mathbb{S}^2$ for parametrization here. These two shapes are naturally similar, but still there are some subtle differences of features between them. When we perform the optimization over Γ , the cost function reduces from 0.047 to 0.0311, a significant decrease although it is hard to visualize it on the surface $(f_2 \circ \gamma^*)$. The deformation field $\gamma^*(s) - s$ and its magnitude show the areas where the mesh

has deformed the most during this matching.

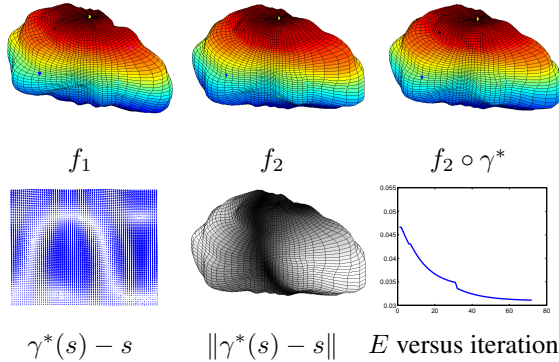


Figure 8. Top: Initial and final matching between pallidum surfaces for subject 1 and subject 2. Bottom: Deformation and cost function. $d([q_1], [q_2]) = 0.0311$

Next, we compare two different anatomical structures for the same subject, the left putamen and the left thalamus. Once again we use $D = S^2$. The putamen and the thalamus surfaces differ much more than two inter-subject pallidum surfaces. For this reason, we expect the distance here to be significantly greater than the distance in the previous example. In fact, this is the case. We observe a decrease in the cost function from approximately 0.117 to 0.0868 due to optimization over Γ , which is a 26% decrease. The deformation fields once again show the areas of the grid where the most deformation occurred.

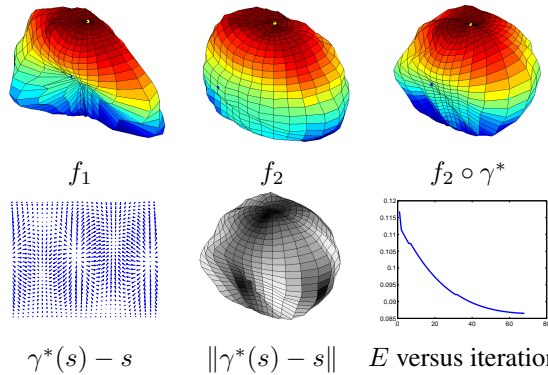


Figure 9. Top: Initial and final matching between a putamen and a thalamus. Bottom: Deformation and cost function. $d([q_1], [q_2]) = 0.0868$

The final example we present is for two hemispherical surfaces. In this case, we would like to match two cropped facial surfaces of the same person, one with a neutral expression and one with an open mouth. We obtain a 19% decrease in the distance function due to optimization over Γ . In addition, we observe a very intuitive result. The deformation fields show the most grid changes in the area of

the lips, which is where the two cropped facial surfaces differ most. By deforming the grid there, we are able to better match the two surfaces.

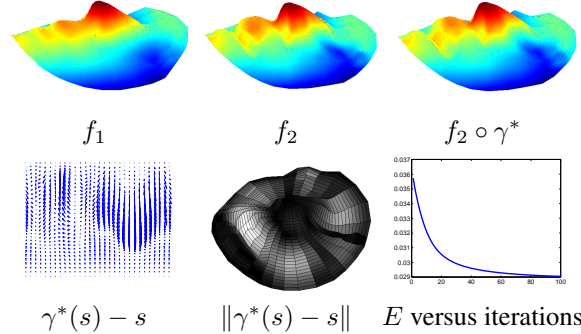


Figure 10. Top: Initial and final matching between face 1 and face 2. Bottom: Deformation and cost function. $d([q_1], [q_2]) = 0.0290$

4.2. Shape-Based Clustering and Classification

In this section we consider clustering and classification of some anatomical structures and cropped facial surfaces using the Riemannian distances. In the first example, we take 10 subjects and four anatomical parts (left putamen, left caudate, left pallidum, and left thalamus) for each subject. An example of each of these structures, for the same subject, is shown in the top part of Figure 11. For these 40 objects, we compute pairwise Riemannian distances using Algorithm 1. Using this distance matrix, and a simple multidimensional scaling algorithm, we plot these individual objects in \mathbb{R}^3 using the first three MDS Euclidean coordinates in the bottom panel. It can be seen that this distance function successfully separates the four anatomical parts.

We also use Riemannian distances between left putamens of young adults who were cases and controls in an attention deficit hyperactivity disorder (ADHD) study. These datasets were selected from the Detroit Prenatal Alcohol and Cocaine Exposure Cohort. In these datasets, 5 subjects were diagnosed with ADHD (labeled 1,3,4,6,10) and 5 subjects were healthy (labeled 2,5,7,8,9). Figure 12 displays the disease classification results. We achieve an 80% disease classification rate using the leave one out nearest neighbor classifier. The right MDS plot shows a clear division between the case and control groups. Subjects 8 and 10 are significantly different from everyone else, and therefore, they cluster on their own.

Another example involves classification of cropped facial surfaces. We perform nearest neighbor classification on a data set of cropped facial surfaces for 25 people. For each person, we have two neutral facial expressions and one non-neutral expression (e.g. smile, open mouth). We use one of the neutral expressions for each person as the gallery

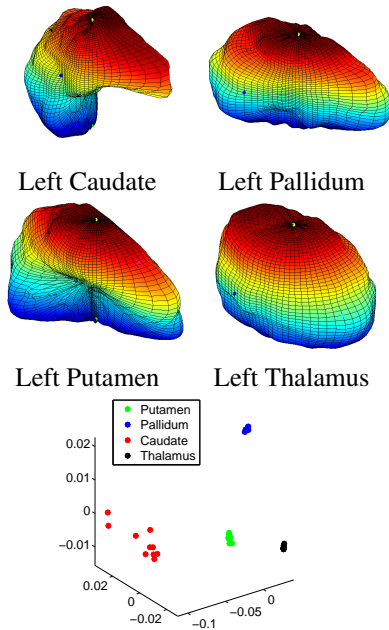


Figure 11. Clustering of anatomical parts (putamen, caudate, pallidum, and thalamus) using shape metrics.




Figure 12. Left: Dendrogram display of clustering results. Right: Nonmetric multidimensional scaling display of clustering results. Blue (small)=Controls; Red (large)=Cases

and the remaining neutral expressions and the non-neutral expressions as the probe. The results of this experiment are summarized in Figure 13. On the left, we display the 25x50 distance matrix. The darkest regions in the distance matrix occur on the diagonal, which shows that the distance between cropped facial surfaces for the same person are very small. On the right, we present the plot of recognition rate versus the rank of the nearest neighbor classifier. It is evident that our methodology has some success in this application. When the rank of the nearest neighbor classifier is 1, we achieve a 68% classification rate. When we increase the classifier rank to 5, we see an increase in the classification rate to 88%.

5. Conclusion

Shape analysis of 3D objects is very important in many scientific fields. We have proposed a novel approach for the shape analysis of 3D objects, which is invariant to rigid motion, scaling and most importantly re-parametrization of surfaces. In order to remove the re-parametrization group from the representation space, we have utilized a gradient




Figure 13. Left: Matrix of Riemannian distances between cropped facial surfaces. Right: Diagram of classification rate versus classifier rank for faces of 25 different people.

descent algorithm. In addition, we have defined a proper Riemannian distance between surfaces, which may be used for classification purposes. Finally, we have presented some examples of this framework being applied to shape matching of toy surfaces, anatomical surfaces, and cropped facial surfaces. We have not presented any results of geodesic paths. We plan to address the issue of geodesics between surfaces in a future paper.

Acknowledgements

This research was supported in part by AFOSR FA9550-06-1-0324, ONR N00014-09-1-0664 and NSF DMS-0915003.

References

- [1] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. 1
- [2] S. Bouix, J. C. Pruessner, D. L. Collins, and K. Siddiqi. Hippocampal shape analysis using medial surfaces. *NEUROIMAGE*, 25:1077–1089, 2001. 1
- [3] I. L. Dryden and K. Mardia. *Statistical Shape Analysis*. John Wiley & Son, 1998. 1
- [4] M. Fuchs, B. Jüttler, O. Scherzer, and H. Yang. Shape metrics based on elastic deformations. *Journal of Mathematical Imaging and Vision*, 35(1):86–102, 2009. 1
- [5] U. Grenander and M. I. Miller. Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics*, LVI(4):617–694, 1998. 1
- [6] X. Gu and B. C. Vemuri. Matching 3d shapes using 2d conformal representations. In *Proc. of MICCAI, Volume 3216*, pages 771–780, 2004. 3
- [7] S. H. Joshi, E. Klassen, A. Srivastava, and I. H. Jermyn. A novel representation for riemannian analysis of elastic curves in \mathbb{R}^n . In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–7, 2007. 1, 2
- [8] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. In *Proc. of SIGGRAPH*, 2006. 2
- [9] S. Osher and R. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer Verlag, 2003. 1
- [10] M. Styner, I. Oguz, S. Xu, C. Brechbuhler, D. Pantazis, J. Levitt, M. Shenton, and G. Gerig. Framework for the statistical shape analysis of brain structures using spharm-pdm. In *MICCAI Open Science Workshop*, 2006. 2
- [11] L. Younes, P. W. Michor, J. Shah, D. Mumford, and R. Lincei. A metric on shape space with explicit geodesics. *Matematica E Applicazioni*, 19(1):25–57, 2008. 1