

Elastic Shape Matching of Parameterized Surfaces Using Square Root Normal Fields

Ian H. Jermyn¹, Sebastian Kurtek², Eric Klassen³, and Anuj Srivastava⁴

¹ Department of Mathematical Sciences, Durham University, Durham, England

² Department of Statistics, The Ohio State University, Columbus, Ohio

³ Department of Mathematics, Florida State University, Tallahassee, Florida

⁴ Department of Statistics, Florida State University, Tallahassee, Florida

Abstract. In this paper we define a new methodology for shape analysis of parameterized surfaces, where the main issues are: (1) choice of metric for shape comparisons and (2) invariance to reparameterization. We begin by defining a general elastic metric on the space of parameterized surfaces. The main advantages of this metric are twofold. First, it provides a natural interpretation of elastic shape deformations that are being quantified. Second, this metric is invariant under the action of the reparameterization group. We also introduce a novel representation of surfaces termed square root normal fields or SRNFs. This representation is convenient for shape analysis because, under this representation, a reduced version of the general elastic metric becomes the simple \mathbb{L}^2 metric. Thus, this transformation greatly simplifies the implementation of our framework. We validate our approach using multiple shape analysis examples for quadrilateral and spherical surfaces. We also compare the current results with those of Kurtek et al. [1]. We show that the proposed method results in more natural shape matchings, and furthermore, has some theoretical advantages over previous methods.

1 Introduction

Shape analysis plays a very important role in image analysis and computer vision, with applications in medical diagnostics, bioinformatics, graphics, robotics, target recognition, and many more. There has been an increasing interest in using a Riemannian framework for shape analysis of objects due to the breadth of tools that it offers. First, it allows us to remove all shape-preserving transformations from the representation space, using the notion of quotient spaces. Second, 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 [2, 3]. To the best of our knowledge, there are very few papers that study shapes of continuous 2D surfaces in a similar fashion.

There have been many different representations used to study shapes of 3D objects. Several groups have proposed methods based on deformable templates where the surfaces are embedded in 3D domains [4]. While these methods are both prominent and pioneering in medical image analysis, they are typically computationally expensive since the registration is performed on the full volumes.

An alternative approach is based on manually-generated landmarks under the Kendall shape theory [5]. Others study 3D shape variabilities using level sets [6], curvature flows [7], or point cloud matching via the iterative closest point algorithm [8]. Also, there has been remarkable success in the use of medial representations for shape analysis, especially in medical image analysis [9, 10].

However, the most natural representation for studying shapes of 3D objects seems to be based on their continuous boundaries, but the parameterization variability makes it difficult. Some papers use SPHARM or SPHARM-PDM [11, 12] to tackle this problem by choosing a fixed arc-length type parameterization. This is a major restriction and does not result in elastic shape analysis of surfaces. A large set of papers in the literature treat the parameterization (or registration) and comparison steps in a disjoint manner [13–15]. In other words, they take a set of surfaces and use some energy function, such as the entropy or the minimum description length to register points across surfaces. Once the surfaces are registered, they are compared using standard procedures. Because these two steps are often performed under different metrics, the resulting registrations and shape comparisons tend to be suboptimal.

A recent paper by Fuchs et al. [16] considers both the curves and their interiors, but falls short of studying 3D objects. Windheuser et al. [17] solve a dense registration problem but use a linear interpolation between registered pairs of points in \mathbb{R}^3 to compute geodesic paths. Another paper by Kilian et al. [18] represents parameterized surfaces by discrete triangulated meshes, assumes a Riemannian metric on the space of such meshes, and computes geodesic paths between given meshes. 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. In contrast, we would like to remove the reparameterization variability (akin to performing registration of surfaces) so that different surfaces with the same shape but different parameterizations have zero distance between them.

In order to tackle this problem, it is helpful to first look at how these goals were achieved in shape analysis of parameterized curves. In many papers the authors define a new representation of curves such that under the standard \mathbb{L}^2 metric the reparameterization of curves preserves distances. Let $\beta_1, \beta_2 : [0, 1] \rightarrow \mathbb{R}^n$ be two parameterized curves, and let Γ be the reparameterization group (the set of all 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. Srivastava et al. [3] define a function $q(t) = \dot{\beta}(t)/|\dot{\beta}(t)|^{\frac{1}{2}}$ that is used to represent a curve β , with the following desirable properties. First, since the q -function for a re-parameterized curve $\beta \circ \gamma$ is given by $(q, \gamma) = (q \circ \gamma)\sqrt{\dot{\gamma}}$, we have that $\|q_1 - q_2\| = \|(q_1, \gamma) - (q_2, \gamma)\|$ for all $\gamma \in \Gamma$. In this way, one can remove the reparameterization group and define a distance between shapes of curves as $\inf_{\gamma \in \Gamma} \|q_1 - (q_2, \gamma)\|$. Second, this \mathbb{L}^2 metric is in fact a specific instance of a family of elastic metrics on the space of curves. This relationship provides useful interpretations of the types of deformations measured by the metric, and justifies its use. Third, while for $n = 2$ there is a one-parameter family of such functions, for $n > 2$, q is unique.

Kurtek et al. [1, 19–21] presented a similar framework for shape analysis of parameterized surfaces. They define a new representation of surfaces termed a q -map such that the \mathbb{L}^2 distance between two q -maps is preserved under the action of the same γ . They are further able to define geodesics and statistics on the shape space of parameterized surfaces. Although this work is a very important step in the right direction, the main drawback of this method is that the q -map representation of surfaces does not have a clear relationship to an underlying elastic metric. The representation was solely devised for convenience of being able to compare the shapes of parameterized surfaces using the \mathbb{L}^2 metric. Another flaw of the q -map representation is that translating two shapes equally does not preserve the distance between them, which can cause major problems for shape analysis.

In this paper we present a representation of surfaces fully analogous to the q -representation of curves. It allows invariant shape analysis of surfaces, and the \mathbb{L}^2 metric in this representation is a special case of a general elastic metric on surfaces. Thus, the main contributions of this paper are the following:

1. The definition of a general elastic metric on the space of parameterized surfaces.
2. The introduction of a novel representation for shape analysis of surfaces based on this metric.

The rest of the paper is as follows. Section 2 introduces the new, elastic metric and representation, and shows how to compute distances between shapes in this metric. Section 3 illustrates the ideas using toy examples as well as more complicated surfaces. We also provide results for image matching using surface graphs. Section 4 concludes.

2 Mathematical Framework

Let \mathcal{F} be the space of all smooth embeddings $f : D \rightarrow \mathbb{R}^3$, where D will be the square or the sphere. Each such embedding defines a surface $f(D) \subset \mathbb{R}^3$. Let Γ be the set of all diffeomorphisms of D . It acts naturally on \mathcal{F} by composition: $(f, \gamma) \mapsto f \circ \gamma$. The map from embeddings to surfaces is not injective: two embeddings related by the action of Γ correspond to the same surface, so that the space of surfaces can be thought of as the quotient \mathcal{F}/Γ .

The tangent space at $f \in \mathcal{F}$, $T_f(\mathcal{F})$, can be identified with \mathcal{F} . Given a fixed measure ds on D , we can define a Riemannian structure on \mathcal{F} : for $v_1, v_2 \in T_f(\mathcal{F})$, the inner product $\langle v_1, v_2 \rangle = \int_D \langle v_1(s), v_2(s) \rangle ds$. The resulting \mathbb{L}^2 distance between $f_1, f_2 \in \mathcal{F}$ is $(\int_D |f_1(s) - f_2(s)|^2 ds)^{\frac{1}{2}}$. While simple, this metric has a critical defect, just as in the case of curves: the action of Γ does not preserve distances. We have

$$d^2(f_1 \circ \gamma, f_2 \circ \gamma) = \int_D |f_1(\gamma(s)) - f_2(\gamma(s))|^2 ds = \int_D |f_1(\tilde{s}) - f_2(\tilde{s})|^2 J_\gamma(\tilde{s})^{-1} d\tilde{s}$$

where J_γ is the Jacobian of γ . This is not equal to $d^2(f_1, f_2)$ in general because the pushforward measure, $J_\gamma^{-1} ds$, is not equal to ds . This lack of isometry means

that the space of surfaces \mathcal{F}/Γ does not inherit the structure of a Riemannian manifold from \mathcal{F} , thereby making this metric difficult to use for analyzing the shapes of surfaces. One solution is to restrict attention to diffeomorphisms that preserve ds [22], but there is no reason in general to impose this restriction on matchings between surfaces. Another solution, proposed by Kurtek et al. [1], is to develop a new representation of surfaces such that the action of Γ preserves \mathbb{L}^2 distances. Although convenient, the representation defined in that work suffered from some fundamental issues described in the previous section. We suggest an alternative approach based on a new elastic Riemannian metric and a convenient representation of surfaces.

2.1 The Elastic Metric for Surfaces

Let $n(s) \in \mathbb{R}^3$ denote the normal vector to the surface at the point $s \in D$. Let $(u, v) : D \rightarrow \mathbb{R}^2$ be coordinates on (a chart of) D ; then $n(s) = \frac{\partial f}{\partial u}(s) \times \frac{\partial f}{\partial v}(s)$. The area measure induced on D by f at s is given by $r(s) = |n(s)|$, so the normalized normal vector is $\tilde{n}(s) = \frac{n(s)}{r(s)}$. We will represent the embedding f using the pair (r, \tilde{n}) . Let $(\delta r_1, \delta \tilde{n}_1)$ and $(\delta r_2, \delta \tilde{n}_2)$ be two tangent vectors to the representation space at (r, \tilde{n}) . The new metric is then

$$\langle (\delta r_1, \delta \tilde{n}_1), (\delta r_2, \delta \tilde{n}_2) \rangle_{(r, \tilde{n})} = \frac{1}{4} \int_D \frac{\delta r_1(s) \delta r_2(s)}{r(s)} ds + \int_D \langle \delta \tilde{n}_1(s), \delta \tilde{n}_2(s) \rangle r(s) ds. \tag{1}$$

To understand this metric, consider the effect of a small change in f on an infinitesimal patch in the surface around the point $f(s)$. The effect can be decomposed into a change in the normal direction $\tilde{n}(s)$ of the patch and a change in its geometry. The latter can be further decomposed into a change in the area of the patch (*i.e.* in $r(s)$), and an area-preserving change in its shape [23, 24]. The above metric measures the first two types of changes, but does not ‘notice’ changes in f that change the shape of the patch while preserving its area and normal direction.

Eqn. (1) is in fact a special case of a general elastic metric for surfaces that measures all three types of change. Instead of mapping a surface f to its unit normal field and area factor, we can map it to its unit normal field and its full pullback metric, $g = f^*h$, where h is 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 determinant). The general elastic metric for surfaces is then

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

where, for positivity, $\lambda > \frac{-2}{n}$ and $c \in \mathbb{R}_+$. Let us say a few words about this metric:

- Unlike the (r, \tilde{n}) representation, the (g, \tilde{n}) representation is injective (up to isometries of \mathbb{R}^3) [25]. Like the (r, \tilde{n}) representation, it is not surjective.

- The first two terms form the unique family of ultralocal metrics on the space of Riemannian metrics [23]. The $\lambda = 0$ case of this metric is studied in *e.g.* [24, 26, 27].
- The term multiplied by λ can be written as

$$\int_D ds |g|^{\frac{1}{2}} \text{tr}(g^{-1} \delta g_1) \text{tr}(g^{-1} \delta g_2) = 4 \int_D ds |g|^{-\frac{1}{2}} \delta_1(|g|^{\frac{1}{2}}) \delta_2(|g|^{\frac{1}{2}}). \tag{3}$$

Since $r = |g|^{\frac{1}{2}}$, this is the same as the first term in Eqn. (1), while the third term in Eqn. (2) is the same as the second term in Eqn. (1). Thus Eqn. (1) is a special case of Eqn. (2).

The first two terms of Eqn. (2) measure general ‘stretching’ of the surface, *i.e.* changes in both the area and the shape of local patches, while the third term measures changes in the normal direction, that is, ‘bending’. Thus this metric is fully analogous to the elastic metric for curves.

2.2 Surface Representation

Why, in this paper, do we focus on the metric in Eqn. (1) rather than the more general Eqn. (2)? The reason is that despite its appearance, the metric in Eqn. (1) is *Euclidean*: we can find ‘coordinates’ such that this metric takes on the simple \mathbb{L}^2 form. This means that global distances can be found without solving geodesic equations, as simple \mathbb{L}^2 norms of differences, thereby dramatically simplifying shape analysis. We call this new, convenient representation of surfaces the *square root normal field (SRNF)*:

Definition 1. *The square root normal field (SRNF) $q : D \rightarrow \mathbb{R}^3$ is defined as*

$$q(s) = \sqrt{r(s)} \tilde{n}(s) = \frac{n(s)}{\sqrt{r(s)}} = \frac{n(s)}{|n(s)|^{\frac{1}{2}}}. \tag{4}$$

Since $|q|^2 = r$, the \mathbb{L}^2 norm of q is just the area of the surface. Thus the space of SRNFs is a subset of $\mathbb{L}^2(D, \mathbb{R}^3)$, henceforth abbreviated to \mathbb{L}^2 . The map $\mathcal{F} \rightarrow \mathbb{L}^2$ will be denoted φ , with $\mathcal{Q} = \varphi(\mathcal{F})$.

We will now see that the \mathbb{L}^2 metric in the SRNF representation is Eqn. (1). A tangent vector to an SRNF is given by $\delta q(s) = \frac{1}{2\sqrt{r(s)}} \delta r(s) \tilde{n}(s) + \sqrt{r(s)} \delta \tilde{n}(s)$.

Taking the \mathbb{L}^2 inner product between two such vectors, we obtain $\langle \delta q_1, \delta q_2 \rangle = \frac{1}{4} \int_D \delta r_1(s) \delta r_2(s) \frac{1}{r(s)} ds + \int_D r(s) \langle \delta \tilde{n}_1(s), \delta \tilde{n}_2(s) \rangle ds$, since $\langle \tilde{n}(s), \delta \tilde{n}_i(s) \rangle = 0$.

Our approach to shape analysis of surfaces will thus be to represent surfaces using their SRNFs, and use the \mathbb{L}^2 metric to compare their shapes.

2.3 Group Actions, Shape Invariances, and Distance

Our goal is to define a distance between *shapes*, *i.e.* surfaces modulo translation, rotations, and scaling. To do this, we have to define a distance, not between elements of \mathcal{F} , but between equivalence classes defined by the action, on \mathcal{F} , of Γ and the similarity group of \mathbb{R}^3 . If the SRNF is to be useful, these groups must also act in a consistent way on \mathbb{L}^2 . The following facts are easily checked:

- The SRNF is invariant to translations since it involves only derivatives of f .
- Scaling the surface $f \mapsto af$ sends $q \mapsto \sqrt{a}q$.
- Rotating the surface $f \mapsto Of$ sends $q \mapsto Oq$.
- The action of Γ sends $q \mapsto (q, \gamma) = \sqrt{J_\gamma}(q \circ \gamma)$. We extend this to all \mathbb{L}^2 .
- The actions of $SO(3)$ and Γ commute on both \mathcal{F} and \mathbb{L}^2 , meaning that there is a well-defined action of the product group $\mathcal{G} = SO(3) \times \Gamma$ on both these spaces. Let $\mathcal{S} = \mathcal{F}/\mathcal{G}$ and $\mathcal{R} = \mathbb{L}^2/\mathcal{G}$.
- The map φ is equivariant with respect to the action of \mathcal{G} on \mathcal{F} and \mathbb{L}^2 . Thus there exists a well-defined map $[\varphi] : \mathcal{S} \rightarrow \mathcal{R}$.
- The action of \mathcal{G} on \mathbb{L}^2 is by isometries: $\|q_1 - q_2\| = \|(Oq_1, \gamma) - (Oq_2, \gamma)\|$.

The first point guarantees that shape analysis will be invariant to translations. To produce invariance to scaling, we scale all surfaces to unit area. Note that this amounts to restricting attention to the unit sphere in \mathbb{L}^2 . We then define a distance on \mathcal{R} by minimizing over the equivalence classes of \mathcal{G} :

Definition 2

$$d([q_1], [q_2]) = \inf_{(O, \gamma) \in \mathcal{G}} \|q_1 - (Oq_2, \gamma)\| . \tag{5}$$

We can now define a distance between shapes $[f_1]$ and $[f_2]$ by $d_{\mathcal{S}}([f_1], [f_2]) = d([\varphi]([f_1]), [\varphi]([f_2]))$. Note that d is an *extrinsic* metric, in two senses. First, because unit-area surfaces live on the unit sphere in \mathbb{L}^2 , and the above is a chordal distance. Second, because it is not clear that there will always be a path in \mathcal{R} between $[q_1]$ and $[q_2]$ with length $d([q_1], [q_2])$ that lies entirely in the image $[\varphi](\mathcal{S})$ of $[\varphi]$. Note too that because φ is not injective, $d_{\mathcal{S}}$ is a pseudometric on \mathcal{S} . We discuss this point further in section 4.

2.4 Gradient over Reparameterization Group

The computation of d in Eqn. (5) requires solving the joint optimization problem on $\mathcal{G} = SO(3) \times \Gamma$. For a fixed $\gamma \in \Gamma$, the minimization over $SO(3)$ can be performed directly using 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 singular value decomposition $A = U\Sigma V^T$, we can define the optimal rotation as $O^* = UV^T$ (if the determinant of A is negative, the last column of V changes sign). Thus if we can optimize over Γ for fixed O , then we can alternate between minimizing over $SO(3)$ and Γ , and thereby converge to a solution.

To optimize over Γ , we use a gradient descent approach. This is very similar to the approach presented in [1, 19], because the action of Γ is the same in the two cases. Nonetheless, we sketch out an outline of this procedure here for convenience.

At the current iteration, let the parameterization of the second surface result in $q_2 \in [q_2]$. Let $T_{q_2}([q_2])$ be the space of tangents to the Γ orbit $[q_2]$ at q_2 and set $v = \frac{q_1 - q_2}{|q_1 - q_2|}$ in $T_{q_2}([q_2])$. Actually v is the gradient of the cost function $E \equiv \|q_1 - q_2\|$, given in Eqn. (5), with respect to q_2 . If the projection of v along

the orbit $[q_2]$ is zero, then the straight line from q_1 to q_2 is perpendicular to the orbit $[q_2]$, and indeed all other orbits; this is the extrinsic geodesic distance between the elements of $[q_1]$ and $[q_2]$. Otherwise, we update q_2 in the direction of this projection. An integral step here is to compute the projection of v on $T_{q_2}([q_2])$. This step is performed numerically: we determine an orthonormal basis of a finite-dimensional subspace of $T_{q_2}([q_2])$ and use those basis elements to approximate the projection of v . These basis elements, in turn, are first derived in the space $T_{\gamma_{id}}(T)$ and then transferred to $T_{q_2}([q_2])$ using an appropriate mapping.

3 Experimental Results

In this section, we present various surface shape comparison results in order to validate the proposed method. We provide examples for both quadrilateral and spherical surfaces. In some cases, we compare the proposed method to the one described in [1]. We also consider the problem of matching and classifying images represented using surface graphs.

3.1 Quadrilateral and Spherical Surface Comparisons

Next, we display examples of matching and comparing quadrilateral and closed surfaces. In each example we display the surfaces with their given parameterizations as f_1 and f_2 . We also provide the optimally registered second surface as $O^*(f_2 \circ \gamma^*)$, the optimal parameterization of the second surface γ^* and the energy at each iteration. The matching of points across surfaces is displayed by mapping the color scheme on surface f_1 to the corresponding points on the second surface. Thus, similar features on both surfaces should be shaded with the same colors. In each example, we also evaluate the obtained registration of surfaces by considering a linear interpolation between them. Note that this interpolation is computed in \mathcal{F} and thus does not represent a geodesic under our metric. Improved registration corresponds to better preservation of features along this path.

Ex. 1 (Fig. 1): Here, we consider an example where the two given surfaces are within a reparameterization of each other. That is, $f_2 = f_1 \circ \gamma$. Thus, we expect the resulting energy to be very close to zero and the linear interpolation between f_1 and $f_2 \circ \gamma^*$ to be a constant path. This is in fact the case.

Ex. 2 (Fig. 1): In this example, we are interested in comparing a surface with one peak to a surface with one high peak and one low peak. We note that, as given, the two surfaces are very misaligned, which can be seen in the linear interpolation between them. After applying the proposed method, we note that the two high peaks match well and that the linear interpolation between the registered surfaces is more natural.

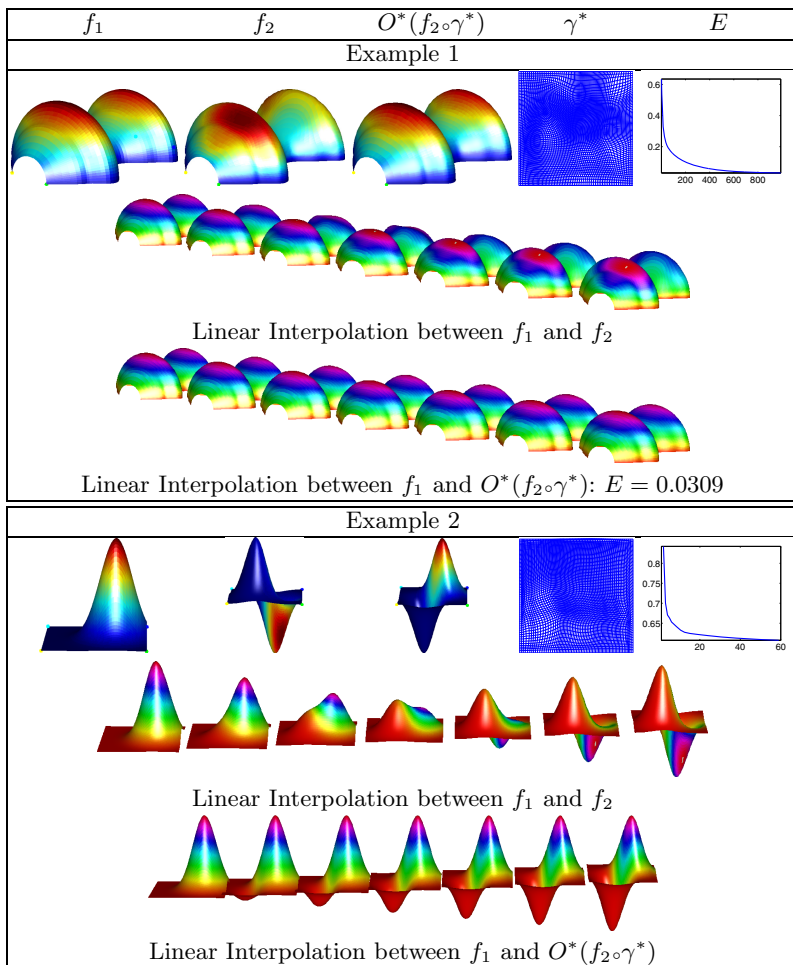


Fig. 1. Comparisons of quadrilateral surfaces. Top: Example for surfaces of revolution where $f_2 = f_1 \circ \gamma$. We expect the linear interpolation to be a constant path and the distance to be close to zero. Bottom: Example for two toy surfaces.

Ex. 3 and 4 (Fig. 2): We present similar results for toy spherical surfaces. Again, we first consider the case where $f_2 = f_1 \circ \gamma$. We note that the proposed method recovers the correct parameterization of the second surface, which is reflected in an energy close to zero and a constant linear interpolation between the surfaces. In Ex. 4, we consider two surfaces with different numbers of peaks. With the given parameterizations, the peaks on the surfaces do not match well, and the surfaces along the linear interpolation have distorted features. But, after applying the proposed method this matching is improved, which again results in a much more natural path between the two surfaces.

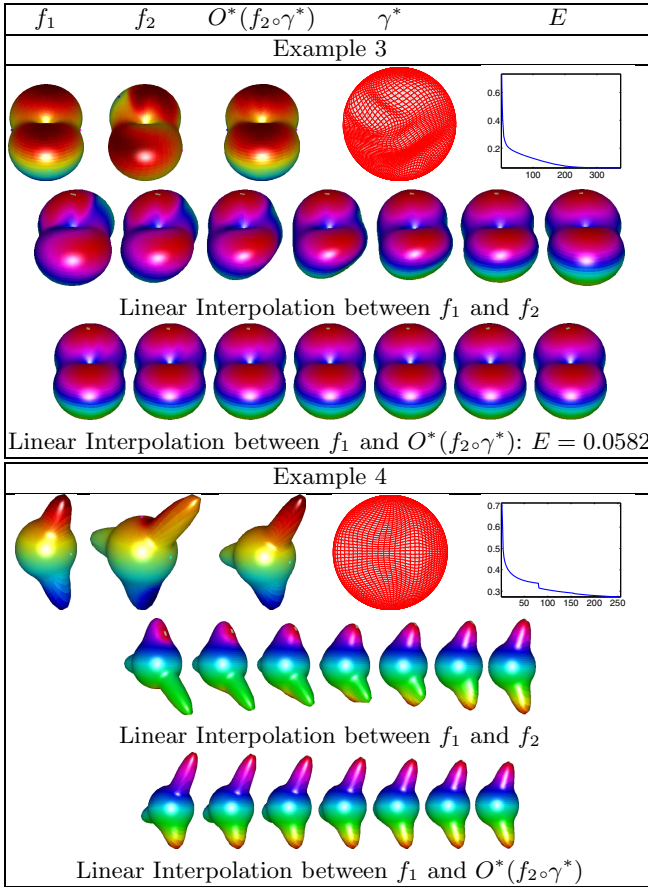


Fig. 2. Comparisons of spherical surfaces. Top: Example where $f_2 = f_1 \circ \gamma$. We expect the linear interpolation to be a constant path and the distance to be close zero. Bottom: Example for two toy surfaces.

Ex. 5 (Fig. 3): Next, we compare the surface matching based on the proposed method to that presented by Kurtek et al. [1]. In this example, the proposed method performs very well at matching the three peaks across the two given quadrilateral surfaces. The latter method attempts to match all of the peaks on f_1 to the largest peak on f_2 , which results in a less natural linear path between the two surfaces.

Ex. 6 (Fig. 4): In this example, we consider two spherical models of a hand. Again, we compare the results obtained using the proposed method and the method presented in [1]. The proposed method yields a much better matching of features across surfaces in this case. Using the registration based on the proposed

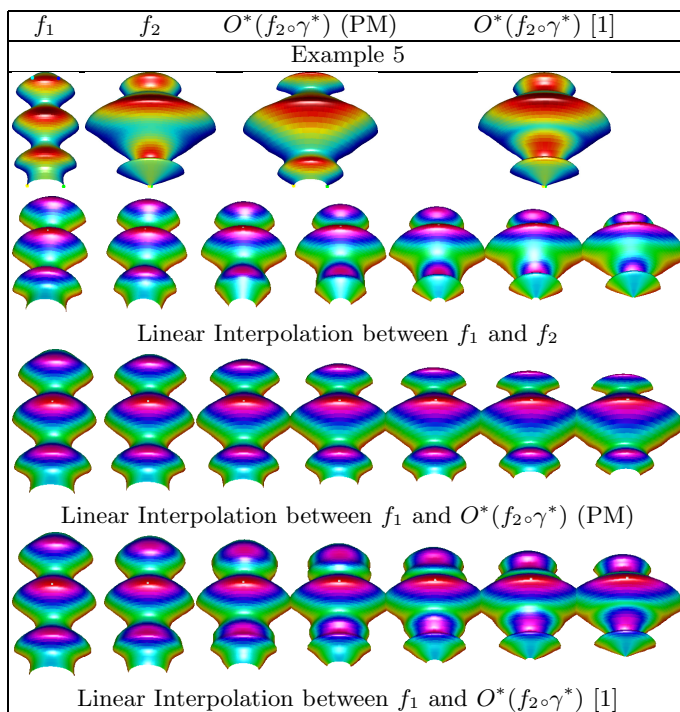


Fig. 3. Surface comparisons based on surfaces of revolution. PM=proposed method.

method the missing finger nicely grows along the linear interpolation. The other method does not provide such an intuitive result.

Ex. 7-12 (Fig. 5): Here, we present more examples of matching some complicated spherical surfaces using the proposed method. We note that the resulting registration provides a natural deformation between the surfaces.

3.2 Image Registration and Classification Using Surface Graphs

In this section, we present registration and classification results for digit images by viewing them as quadrilateral surface graphs. This data is a subset of 100 images (ten for each digit) obtained from the MNIST database.

We begin by showing three examples of matching images with different digits (Fig. 6): two zeros, a one and a zero, and a three and an eight. Again, we compare our method to the one proposed by Kurtek et al. [1]. In each example, we provide the two original images (I_1 and I_2), the optimally warped second image ($I_2 \circ \gamma^*$), and the absolute difference image. We note that in the first example the two methods produce similar results. In the second and third examples, the two

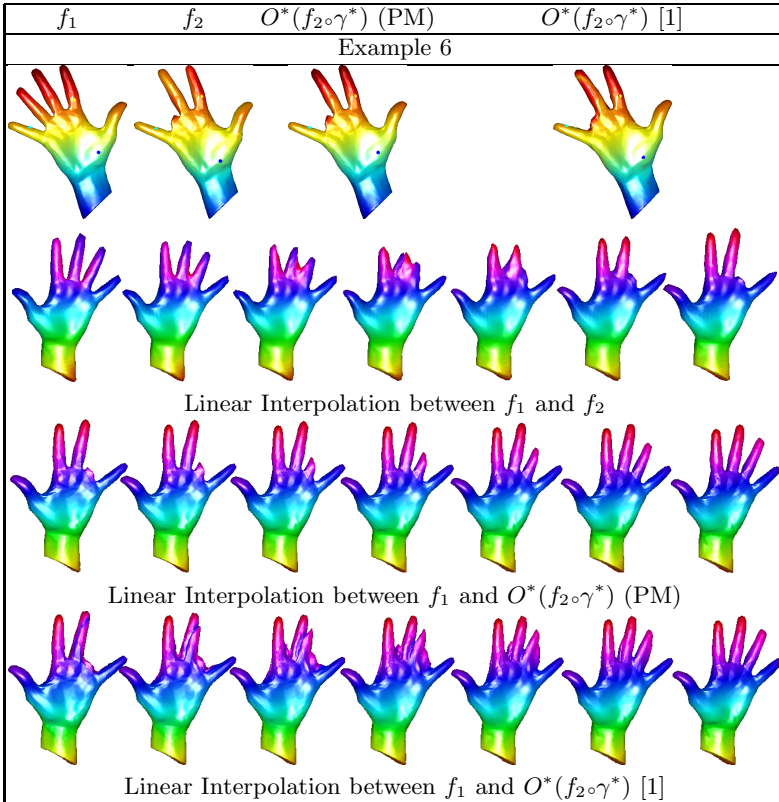


Fig. 4. Surface comparisons based on spherical surfaces

methods provide very different results. The proposed method provides smaller absolute differences between registered images, which is especially evident in the third example.

In order to evaluate the overall quality of this method compared to the one proposed by Kurtek et al., we performed a classification experiment on the 100 digit images (ten zeros, ten ones, ..., ten nines). We began by computing the full pair-wise distance matrices based on the two methods. Using these distance matrices we computed the cluster purity measure and the one, three, and five nearest neighbor classification rates. The results, along with the distance matrices, are reported in Fig. 7. We note that the distance matrix based on the proposed method provides a better separation of classes. This is reflected in a significantly higher cluster purity measure and nearest neighbor classification rates. Furthermore, the proposed method provides more stable classification rates with respect to the number of nearest neighbors considered.

Ex.	f_1	f_2	$O^*(f_2 \circ \gamma^*)$	Lin. Interp.
7				
8				
9				
10				
11				
12				

Fig. 5. More shape matching examples of spherical surfaces




Fig. 6. Image registration using surface graphs. Top: Two zero digits. Middle: A zero and a one. Bottom: A three and an eight.



PM	[1]
 <p>$CP = 0.80$</p> <p>1NN $CR = 0.91$</p> <p>3NN $CR = 0.90$</p> <p>5NN $CR = 0.89$</p>	 <p>$CP = 0.53$</p> <p>1NN $CR = 0.81$</p> <p>3NN $CR = 0.75$</p> <p>5NN $CR = 0.63$</p>

Fig. 7. Digit image classification based on the proposed method and the method defined by Kurtek et al. [1]. CP = cluster purity, NN CR = nearest neighbor classification rate.

4 Summary

We have defined a new elastic metric for shape comparisons of parameterized surfaces. In order to compute distances under this metric, we have introduced a novel representation of surfaces termed square root normal fields. The advantage of using this representation is that the complicated elastic metric simplifies to the standard L^2 metric. We have included a number of results that validate our methodology.

Many questions remain unanswered. The SRNF representation is not faithful, so that many embeddings correspond to the same SRNF. What is the nature of this freedom? Is it useful? Alternatively, we could move to the full elastic metric, Eqn. (2) and the faithful metric/normal field representation. Can we find simplifying coordinates for this metric, analogous to the SRNF? The answers to these questions would further advance the statistical modeling of 3D shape.

References

1. Kurtek, S., Klassen, E., Ding, Z., Srivastava, A.: A novel Riemannian framework for shape analysis of 3D objects. In: Proc. IEEE Computer Vision and Pattern Recognition, pp. 1625–1632 (2010)
2. Younes, L., Michor, P.W., Shah, J., Mumford, D., Lincei, R.: A metric on shape space with explicit geodesics. *Matematica E Applicazioni* 19, 25–57 (2008)
3. Srivastava, A., Klassen, E., Joshi, S., Jermyn, I.: Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans. Pattern Analysis and Machine Intelligence* 33, 1415–1428 (2011)
4. Grenander, U., Miller, M.I.: Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics* LVI, 617–694 (1998)
5. Dryden, I.L., Mardia, K.: *Statistical Shape Analysis*. John Wiley & Son (1998)
6. Malladi, R., Sethian, J., Vemuri, B.: A fast level set based algorithm for topology-independent shape modeling. *Journal of Math. Imag. and Vision* 6, 269–290 (1996)
7. Gu, X., Wang, S., Kim, J., Zeng, Y., Wang, Y., Qin, H., Samaras, D.: Ricci flow for 3D shape analysis. In: Proc. IEEE Int. Conf. on Computer Vision (2007)

8. Almhdie, A., Léger, C., Deriche, M., Lédée, R.: 3D registration using a new implementation of the ICP algorithm based on a comprehensive lookup matrix: Application to medical imaging. *Pattern Recognition Letters* 28, 1523–1533 (2007)
9. Bouix, S., Pruessner, J.C., Collins, D.L., Siddiqi, K.: Hippocampal shape analysis using medial surfaces. *Neuroimage* 25, 1077–1089 (2001)
10. Gorczowski, K., Styner, M., Jeong, J., Marron, J., Piven, J., Hazlett, H., Pizer, S., Gerig, G.: Multi-object analysis of volume, pose, and shape using statistical discrimination. *IEEE Trans. Pattern Analysis and Machine Intelligence* 32, 652–666 (2010)
11. Brechbühler, C., Gerig, G., Kübler, O.: Parameterization of closed surfaces for 3D shape description. *Computer Vision and Image Understanding* 61, 154–170 (1995)
12. Styner, M., Oguz, I., Xu, S., Brechbuhler, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. In: *Proc. MICCAI Open Science Workshop* (2006)
13. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Eurographics State-of-the-Art Report* (2010)
14. Cates, J., Meyer, M., Fletcher, P., Whitaker, R.: Entropy-based particle systems for shape correspondence. In: *Proc. MICCAI Mathematical Foundations of Computational Anatomy*, pp. 90–99 (2006)
15. Davies, R., Twining, C., Cootes, T., Taylor, C.: Building 3-d statistical shape models by direct optimization. *IEEE Trans. Medical Imaging* 29, 961–981 (2010)
16. Fuchs, M., Jüttler, B., Scherzer, O., Yang, H.: Shape metrics based on elastic deformations. *Journal of Mathematical Imaging and Vision* 35, 86–102 (2009)
17. Windheuser, T., Schlickewei, U., Schmidt, F.R., Cremers, D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In: *Proc. International Conference on Computer Vision, ICCV* (2011)
18. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: *Proc. of SIGGRAPH* (2006)
19. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J., Avison, M., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. *IEEE Trans. Medical Imaging* 30, 849–858 (2011)
20. Kurtek, S., Klassen, E., Gore, J., Ding, Z., Srivastava, A.: Elastic geodesic paths in shape space of parametrized surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence* (2010)
21. Kurtek, S., Klassen, E., Ding, Z., Avison, M., Srivastava, A.: Parameterization-Invariant Shape Statistics and Probabilistic Classification of Anatomical Surfaces. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011*. LNCS, vol. 6801, pp. 147–158. Springer, Heidelberg (2011)
22. Gu, X., Vemuri, B.C.: Matching 3D Shapes Using 2D Conformal Representations. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) *MICCAI 2004*. LNCS, vol. 3216, pp. 771–780. Springer, Heidelberg (2004)
23. DeWitt, B.S.: Quantum Theory of Gravity. I. The Canonical Theory. *Physical Review* 160, 1113–1148 (1967)
24. Clarke, B.: The metric geometry of the manifold of Riemannian metrics over a closed manifold. *Calculus of Variations and Partial Diff. Equations* 39 (2010)
25. Abe, K., Erbacher, J.: Isometric immersions with the same Gauss map. *Mathematische Annalen* 215, 197–201 (1975)
26. Ebin, D.: The manifold of Riemannian metrics. *Proc. Symposia in Pure Mathematics (AMS)* 15, 11–40 (1970)
27. Gil-Medrano, O., Michor, P.W.: The Riemannian manifold of all Riemannian metrics. *Quarterly Journal of Mathematics* 42, 183–202 (1991)