

Numerical Inversion of SRNFs for Efficient Elastic Shape Analysis of Star-Shaped Objects

Qian Xie¹, Ian Jermyn², Sebastian Kurtek³, and Anuj Srivastava¹

¹ Florida State University, Tallahassee, Florida, United States
qxie@stat.fsu.edu, anuj@fsu.edu

² Durham University, Durham, County Durham DH1, United Kingdom
ian.jermyn@inria.fr

³ Ohio State University, Columbus, Ohio, United States
kurtek.1@stat.osu.edu

Abstract. The elastic shape analysis of surfaces has proven useful in several application areas, including medical image analysis, vision, and graphics.

This approach is based on defining new mathematical representations of parameterized surfaces, including the square root normal field (SRNF), and then using the \mathbb{L}^2 norm to compare their shapes. Past work is based on using the pull-back of the \mathbb{L}^2 metric to the space of surfaces, performing statistical analysis under this induced Riemannian metric. However, if one can estimate the inverse of the SRNF mapping, even approximately, a very efficient framework results: the surfaces, represented by their SRNFs, can be efficiently analyzed using standard Euclidean tools, and only the final results need be mapped back to the surface space. Here we describe a procedure for inverting SRNF maps of star-shaped surfaces, a special case for which analytic results can be obtained. We test our method via the classification of 34 cases of ADHD (Attention Deficit Hyperactivity Disorder), plus controls, in the Detroit Fetal Alcohol and Drug Exposure Cohort study. We obtain state-of-the-art results.

Keywords: Statistical shape analysis, elastic shape analysis, parameterized surface, geodesic computation, deformation analysis.

1 Introduction

The analysis of the shapes of 3D objects is an important area of research with a wide variety of applications. The need for shape analysis arises in many branches of science, for example, medical image analysis, protein structure analysis, computer graphics, and 3D printing and prototyping. Many of these are especially concerned with capturing variability within and across shape classes, and so the main focus of research has been on statistical shape analysis and on comparing shapes [2,22,28]. The main differences among the different approaches proposed so far lie in the mathematical representations and metrics used in the analysis. One may use chosen landmarks to represent shapes, and perform Kendall-type shape analysis [8], or use point clouds and apply thin plate splines or ICP [3]. One may represent shapes using medial surfaces [4], level sets [21], or deformable templates [11]. However, the most natural representation for studying

the shapes of 3D objects would seem to be their continuous boundaries. Windheuser et al. [26] solve a dense registration problem, but use linear interpolation between registered pairs of points in \mathbb{R}^3 to compute geodesic paths. Kilian et al. [17] represent parameterized surfaces by discrete triangulated meshes, assume a Riemannian metric on the space of such meshes, and compute geodesic paths between given meshes. The method has the limitation that it assumes the correspondence between points on the two meshes to be known. Heeren et al. [12] propose a method to compute geodesic-based deformations of thin shell shapes. Some papers use SPHARM or SPHARM-PDM [5,24] to tackle this problem by choosing a fixed arc-length type parameterization. This is a major restriction, and does not allow elastic shape analysis of surfaces. They also assume that the surfaces are already in full correspondence. A large set of papers in the literature treat the parameterization (or registration) and comparison steps in a disjoint manner [4,29,10,7,25]. 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.

Recently there has been increasing interest in frameworks for studying the shapes of parameterized surfaces, and in particular in methods that provide invariance to shape-preserving transformations such as rigid motions, global scaling, and reparameterizations. These frameworks are predominantly Riemannian: one identifies an appropriate representation space for the relevant surfaces, endows it with a Riemannian metric, and develops an algorithm for computing geodesic paths under that metric. Invariance to shape-preserving transformations is obtained by forming quotient spaces under these groups, and geodesic calculations are then transferred to this quotient space using an *alignment* step. The key idea is to choose a mathematical representation and an associated Riemannian metric so that the desired invariances are obtained, and geodesic computations are rendered simple. This has been achieved in the shape analysis of **curves** by using as representation and metric, the *square-root velocity function* (SRVF) and a particular member of the family of *elastic metrics*: the resulting metric in the SRVF space is then the \mathbb{L}^2 metric [23]. The \mathbb{L}^2 metric greatly simplifies computations, and enables sophisticated statistical analyses that require fast geodesic calculations. Critical to its utility is the fact that the mapping from the space of curves to the SRVF space is a bijection (up to a translation). Solutions found in SRVF space using the \mathbb{L}^2 metric can thus be uniquely mapped back to the original curve space, which is significantly more efficient than calculating in the curve space itself. This paper contributes to the search for a similarly efficient framework for the shape analysis of **surfaces**.

Kurtek et al. [18,20] took the first steps in this direction. Let $f : \mathbf{S}^2 \rightarrow \mathbb{R}^3$ be a parameterized surface and let \mathcal{F} be the space of all such smooth mappings. Suppose \mathbf{S}^2 is parameterized by the pair $s \equiv (u, v)$ for all $s \in \mathbf{S}^2$. Kurtek et al. introduced a surface representation defined by $q(s) = \sqrt{|n(s)|}f(s)$, where $n(s) = f_u(s) \times f_v(s)$ is the unnormalized normal to the surface at $f(s)$; this was termed the *square-root map* (SRM). Equipping the space of SRMs with the \mathbb{L}^2 metric greatly simplifies geodesic calculations, and also has the crucial property that Γ , the group of all orientation-preserving diffeomorphisms of \mathbf{S}^2 , acts by isometries. Unfortunately, the representation

has several limitations, including that the metric distance between two shapes changes if they are both translated by the same amount; that it is difficult to invert (indeed may not have an inverse); and that the metric has no clear physical interpretation in terms of surfaces.

Jermyn et al. [16] introduced a new representation that avoids some of the limitations of the SRM, while preserving its advantages: the *square-root normal field* (SRNF) sends $f \mapsto Q(f)$, where $Q(f)(s) \equiv n(s)/|n(s)|^{1/2}$. Equipping the space of SRNFs, \mathcal{Q} , with the \mathbb{L}^2 metric again trivializes geodesic calculations, and Γ again acts by isometries. Now, however, the representation is translation invariant by definition, while the \mathbb{L}^2 metric on \mathcal{Q} corresponds to a partial elastic Riemannian metric on \mathcal{F} strictly analogous to the elastic metric used in the case of curves.

The SRNF shares one difficulty with the SRM, however, and that is the problem of inversion. Knowing $Q(f)$ is equivalent to knowing the Gauss map $\tilde{n} = n/|n|$ and the induced measure $|n|^{1/2}$. While the Gauss map together with the induced *metric* is sufficient to reconstruct the surface up to translations and rotations [1,9] (or in combination with only the conformal class of the metric, up to translations, rotations and scale [13]), it is not clear that Q is injective up to simple transformations.¹ (In addition, Q is almost certainly not surjective, a point to which we will return.) If one cannot invert the representation, geodesics and statistical analyses conducted in \mathcal{Q} cannot be moved back to \mathcal{F} . One can always pull the \mathbb{L}^2 metric back to \mathcal{F} and perform computations there [27], but this defeats the purpose of introducing the representation and the \mathbb{L}^2 metric on it.

An alternative is to proceed pragmatically, supposing invertibility until it creates problems. (It is worth noting that even if f is not unique given $Q(f)$, the distance between any two such surfaces is zero, and thus any two geodesics in \mathcal{F} mapping to a geodesic in \mathcal{Q} will have the same length.) We take this pragmatic approach in this paper. The problem we wish to solve is this: Given $q \in \mathcal{Q}$, find $f \in \mathcal{F}$ such that $Q(f) = q$. Were it solved, geodesics, mean shapes, PCA, etc. could be computed in \mathcal{Q} under the \mathbb{L}^2 metric and then mapped back to \mathcal{F} , just as is possible in the case of curves using the SRVF, with resulting large gains in computational efficiency with respect to *e.g.* [16,27]. For general surfaces, this can only be done numerically. We develop a numerical method to find such an f if it exists, and to find the closest (in the elastic metric) f to the set $Q^{-1}(q)$ if it does not. This numerical procedure is expensive, however, and in this paper we do not use it directly to invert Q for general surfaces. Rather, we show that for an important subset of surfaces, an *analytic* solution exists to the inversion problem. These are the ‘star-shaped’ surfaces, *i.e.* those whose enclosed volumes are star domains, a large family of surfaces with great relevance for many real problems. Combining the analytic result with the numerical procedure, we are able to compute geodesics and perform statistical analyses in the space of star-shaped surfaces in a very efficient manner: in fact the computational cost is reduced by an order of magnitude.

The paper is organized as follows. Section 2 gives an overview of the statistical tasks we use as points of comparison, and describes algorithms for these tasks under previous and proposed methods. Section 3 describes the analytic solution to the inversion

¹ It is not simply a case of applying Bonnet’s theorem, because in addition to dn , the second fundamental form involves the derivative df , which is the quantity we are trying to find.

problem, while Section 4 describes the algorithms in detail. Section 5 describes the use of the methods for the classification of subjects with Attention Deficit Hyperactivity Disorder (ADHD) using the shapes of brain subcortical structures, and demonstrates state-of-the-art classification results at greatly reduced computational cost.

2 Model Problems

In order to illustrate the advantages of the new methods, we have selected as points of comparison, several algorithmic and computational tasks that are fundamental to statistical shape analysis:

1. **Geodesic Path Construction:** Given two surfaces f_1 and f_2 , one wants to construct a geodesic path $\alpha(t)$ s.t. $\alpha(0) = f_1$ and $\alpha(1) = f_2$.
2. **Shooting Geodesics:** Given a surface f and a tangent vector v_0 at f , one wants to construct a geodesic path $\alpha(t)$ s.t. $\alpha(0) = f$ and $\dot{\alpha}(0) = v$.
3. **Statistical Summaries of Shapes:** Given a sample of observed surfaces f_1, \dots, f_n , one wants to estimate the mean shape and principal directions of variation.
4. **Random Sampling from Shape Models:** Given a sample of observed surfaces f_1, \dots, f_n , one wants to fit a probability model to the data and sample random shapes from it.
5. **Transferring Deformations between Shapes:** Given surfaces f_1, h_1 and f_2 , one wants to find h_2 such that f_2 deforms to it in a similar way f_1 deforms to h_1 .

Table 1 outlines the algorithms for performing these tasks using both previous and the proposed methods. Computationally intensive steps are underlined, and the computational complexity is indicated in boxes.

3 The Inversion Problem

In order to exploit the SRNF to full advantage, we need to be able to find a surface f such that $Q(f) = q$. In this section, we describe solutions to this problem, first for arbitrary surfaces, and then for star-shaped surfaces.

3.1 General Surfaces

We formulate the inversion problem as an optimization problem by defining an energy function $E_0 : \mathcal{F} \rightarrow \mathbb{R}_+$ such that

$$E_0(f; q) = \|Q(f) - q\|_2^2. \quad (1)$$

Finding an $f \in \mathcal{F}$ such that $Q(f) = q$ is then equivalent to seeking zeros of E_0 . If no such f exists, then a minimizer of E will be a nearest such f under the elastic metric. We define $f^* = \arg \min_{f \in \mathcal{F}} E_0(f; q)$.

Minimization is performed using a gradient descent approach. Since \mathcal{F} is an infinite-dimensional vector space, we will approximate the gradient using a finite basis for \mathcal{F} . From a computational point of view, it may be easier to express the deformation of

Table 1. Comparison of Algorithms

	Previous	Proposed
Karcher Mean	<p>Algorithm 1 Let μ_f^0 be an initial estimate. Set $j = 0$.</p> <ol style="list-style-type: none"> 1. Register f_1, \dots, f_n to μ_f^j. 2. For each $i = 1, \dots, n$, construct a geodesic to connect f_i to μ_f^j and evaluate $v_i = \exp_{\mu_f^j}^{-1}(q_i)$. 3. Compute the average direction $\bar{v} = \frac{1}{n} \sum_{i=1}^n v_i$. 4. If $\ \bar{v}\$ is small, stop. Else, update $\mu_f^{j+1} = \exp_{\mu_f^j}(\epsilon \bar{v})$ by shooting a geodesic, $\epsilon \ll 0$, small. 5. Set $j = j + 1$ and return to Step 1. <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">n geodesics per iteration</p>	<p>Algorithm 2 Let $\bar{q} = Q(\mu_f^0)$ with μ_f^0 as an initial estimate. Set $j = 0$.</p> <ol style="list-style-type: none"> 1. Register $Q(f_1), \dots, Q(f_n)$ to \bar{q}. 2. Update the average $\bar{q} = \frac{1}{n} \sum_{i=1}^n q_i$. 3. If change in $\ \bar{q}\$ is small, stop. Else, set $j = j + 1$ and return to Step 1. <p>Find μ_f by <u>inversion</u> s.t. $Q(\mu_f) = \bar{q}$.</p> <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">1 inversion</p>
Parallel Transport	<p>Algorithm 3 Find a geodesic $\alpha(t)$ connecting f_1 to f_2. For $\tau = 1, \dots, m$, do the following.</p> <ol style="list-style-type: none"> 1. Parallel transport $V(\frac{\tau-1}{m})$ from $\alpha(\frac{\tau-1}{m})$ to $\alpha(\frac{\tau}{m})$ and name it $V(\frac{\tau}{m})$. <p>Set $v^{\parallel} = V(1)$.</p> <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">1 geodesic + m parallel transports</p>	<p>Algorithm 4 Parallel transport on \mathbb{L}^2 remains constant.</p> <ol style="list-style-type: none"> 1. Compute $w = Q_{*,f_1}(v)$ (differential of the mapping Q). 2. Find f by <u>inversion</u> s.t. $Q(f) = Q(f_2) + \epsilon w$, ϵ is small. 3. Evaluate $\frac{f-f_2}{\epsilon}$ and set it to be v^{\parallel}. <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">1 inversion</p>
Transfer Deformation	<p>Algorithm 5</p> <ol style="list-style-type: none"> 1. Find a geodesic $\beta(t)$ connecting f_1 to h_1 and evaluate $v = \exp_{f_1}^{-1}(h_1)$. 2. Find a geodesic $\alpha(t)$ connecting f_1 to f_2. Set $V(0) = v$. For $\tau = 1, \dots, m$, do the following. <ol style="list-style-type: none"> (a) Parallel transport $V(\frac{\tau-1}{m})$ from $\alpha(\frac{\tau-1}{m})$ to $\alpha(\frac{\tau}{m})$ and name it $V(\frac{\tau}{m})$. 3. Shoot a geodesic $\beta'(t)$ from f_2 with velocity $v^{\parallel} = V(1)$ and set $h_2 = \beta'(1)$. <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">3 geodesics + m parallel transports</p>	<p>Algorithm 6 Parallel transport on \mathbb{L}^2 remains constant.</p> <ol style="list-style-type: none"> 1. Compute $v = Q(h_1) - Q(f_1)$. 2. Find h_2 by <u>inversion</u> s.t. $Q(h_2) = Q(f_2) + v$. <p style="text-align: center; border: 1px solid black; display: inline-block; padding: 2px;">1 inversion</p>

a surface, rather than the surface itself, using a basis. We therefore set $f = f_0 + w$, where $w = \sum_{b \in \mathcal{B}} \alpha_b b$, with $\alpha_b \in \mathbb{R}$, and where \mathcal{B} forms an orthonormal basis of \mathcal{F} . (In practice, we use spherical harmonics.) Here f_0 denotes the current estimate of f^* , and w is a deformation of f_0 . Then, we minimize the new energy

$$E(w; q) = \|Q(f_0 + w) - q\|_2^2, \tag{2}$$

with respect to w . One can view f_0 as an initial guess of the solution or a known surface with shape similar to the one being estimated. If no initial guess is possible, one can initialize f_0 as a unit sphere or even set $f_0 = 0$.

We need to evaluate the directional derivatives of E . The directional derivative of E at $f_0 + w$ in the direction of b , $\nabla_b E(w; q)$, is given by:

$$\nabla_b E(w; q, f_0) = \frac{d}{d\epsilon} \Big|_{\epsilon=0} \|Q(f_0 + w + \epsilon b) - q\|_2^2 = 2\langle Q(f_0 + w) - q, Q_{*,f_0+w}(b) \rangle. \tag{3}$$

Here $Q_{*,f}$ denotes the differential of Q at f . This can be evaluated using the following expression: for all $s \in \mathbf{S}^2$,

$$Q_{*,f}(b)(s) = \frac{n_b(s)}{\sqrt{|n(s)|}} - \frac{n(s) \cdot n_b(s)}{2|n(s)|^{5/2}} n(s) \tag{4}$$

where $n_b(s) = f_u(s) \times b_v(s) + b_u(s) \times f_v(s)$. From the perspective of numerical accuracy, the second term can be replaced by a more stable form, $\frac{\tilde{n}(s) \cdot n_b(s)}{2\sqrt{|n(s)|}} \tilde{n}(s)$, resulting in

$$Q_{*,f}(b)(s) = \frac{1}{\sqrt{|n(s)|}} \left(n_b(s) - \frac{\tilde{n}(s) \cdot n_b(s)}{2} \tilde{n}(s) \right). \tag{5}$$

Finally, the update is determined by the gradient $\nabla E(f_0; q) = \sum_{b \in \mathcal{B}} (\nabla_b E(b; q, f_0)) b$ obtained using Eqn. 3, 4 and 5.

3.2 Star-Shaped Surfaces

The numerical solution is for general surfaces. However, solving the optimization problem in this general case is difficult due to the high dimensionality of the search space. We now restrict attention to a special subspace of ‘star-shaped’ surfaces. Remarkably, in this case an analytic solution to the inversion problem exists. At the same time, such surfaces are of great relevance for many applications.

By a ‘star-shaped’ surface, we mean a parameterized surface $f \in \mathcal{F}$ that, up to translation, can be written in the form $f(u, v) = r(u, v)e(u, v)$, where $r(u, v) \in \mathbb{R}$, and $e(u, v) \in \mathbf{S}^2$ is the unit vector in \mathbb{R}^3 given in Euclidean coordinates by $e(u, v) = (\cos(u) \sin(v), \sin(u) \sin(v), \cos(v))$. It can be seen by inspection that the form of e means that the angular spherical coordinates (θ, ϕ) of points on the surface are simply given by $(\theta(u, v), \phi(u, v)) = (u, v)$. Note that the volume enclosed by a star-shaped surface is a star domain, that is, there exists a point in the enclosed volume such that the straight line segments from that point to every point on the surface all lie entirely in the enclosed volume, but that in addition to this purely geometric property, we demand that the surface have a particular parameterization.

In the case of star-shaped surfaces, the map Q can be *analytically* inverted, as follows. The radial component of the normal vector n of an star-shaped surface is, by definition, given by

$$n^r(u, v) = \langle n(u, v), e(\theta(u, v), \phi(u, v)) \rangle \tag{6}$$

since $e(\theta, \phi)$ is the radial unit vector in the direction in \mathbb{R}^3 defined by (θ, ϕ) . If the star-shaped surface were in general parametrization, we could not compute n^r because we would not know θ and ϕ , the angular coordinates of the surface we are trying to recover. In the special parameterization, however, the expression just becomes

$$n^r(u, v) = \langle n(u, v), e(u, v) \rangle \tag{7}$$

and this can be calculated. The result is very simple:

$$n^r(u, v) = r^2(u, v) . \tag{8}$$

As a result, given an SRNF q and a parameterization e , the star-shaped surface \tilde{f} corresponding to this q , i.e. such that $Q(\tilde{f}) = q$, takes the form:

$$\tilde{f}(u, v) = \left(\sqrt{|q(u, v)| q^r(u, v)} \right) e(u, v) , \tag{9}$$

where $q^r = \langle q, e \rangle$ is the radial component of q .

Note that \tilde{f} depends on both q and a fixed parameterization $e(u, v)$. If both are known, then Q can be analytically inverted, as above. If a surface encloses a star domain, but is in a general parameterization (and hence not star-shaped by definition), one can still choose to apply Eqn. 9. In this case, the resulting \tilde{f} will not in general be the original surface f , but it may provide a good initialization for solving the reconstruction-by-optimization problem. The numerical inversion method also provides a way to check whether a given SRNF q corresponds to a star-shaped surface: simply construct \tilde{f} and then compute $Q(\tilde{f})$; if one finds $Q(\tilde{f}) = q$, then q corresponds to a star-shaped surface.

One can thus use the analytic result together with numerical inversion to construct geodesics in \mathcal{F} between two star-shaped surfaces. First, find the geodesic in \mathcal{Q} between the corresponding SRNFs, which is trivially a straight line. It is not guaranteed, however, that all intermediate SRNFs correspond to star-shaped surfaces; thus the analytic form \tilde{f} may not be the right inversion. One can use \tilde{f} , however, as an initial guess for the original surface, thereby better initializing the reconstruction-by-optimization problem.

Reconstruction Examples. To explain the inversion problem further, we present results on reconstructing a synthetic surface in Fig. 1. In this experiment, the target surface is f_o which serves as the ground truth. We compute $q_o = Q(f_o)$ and the goal is to recover the target surface f_o with only q_o known. A surface computed using the analytic inversion in Eqn. 9 is shown as \tilde{f} . Using the unit sphere as initialization, the numerical solution to the optimization problem is shown as f^* . In order to check the convergence of the optimization problem, the energy plotted against iterations is shown in the bottom left panel. The energies, $E(\tilde{f}; q_o)$ and $E(f^*; q_o)$ are shown below the respective

surfaces and compared to $E(f_o; q_o) = 0$ if we get a perfect reconstruction. The pixelwise errors, $|\tilde{f}(s) - f_o(s)|$ and $|f^*(s) - f_o(s)|$ are also shown for the analytic and the numerical solutions in that order. The surface from analytic inversion is very close to the targeted ground truth surface with an energy on the order of 10^{-4} ; the numerical method then brings the energy down further towards zero. The two reconstructed surfaces have very small pointwise errors with respect to the ground truth surface.

We also show results on inverting anatomical surfaces in Fig. 2. For these the experiments, all the energies converge to a small value and the constructed surfaces resemble the ground truth surfaces very well.




Fig. 1. Reconstructing a surface from its SRNF. A target surface (f_o) is numerically reconstructed as f^* with initialization as the unit sphere. The energy plot shows the evolution of energies against iterations with initialization as a unit sphere. The analytically inverted surface \tilde{f} is shown for comparison. The energies $E(\tilde{f}; q_o)$ and $E(f^*; q_o)$ are shown correspondingly. The errors between the reconstructed surfaces and the ground truth are shown on the ground truth surface with colors representing the magnitudes, i.e. $|f^*(s) - f_o(s)|$ for all $s \in \mathbf{S}^2$.

4 Statistical Analysis of Surfaces under Inversion

The ability to invert Q enormously simplifies the algorithms used for various analyses. Compared to the previous framework [27], where analysis is performed on a Riemannian manifold, the new framework performs analysis in the \mathbb{L}^2 space of SRNFs, and only brings the results to the shape space at the very end (Fig. 3).

The basic algorithms for computing the Karcher mean shape, for parallel transport, and for transferring deformations from one shape to another are described in Table 1. Here, we elaborate on the list of target analyses and the mechanisms under inversion.

1. **Geodesic Path Reconstruction:** Given two surfaces f_1 and f_2 , one wants to construct a geodesic path $\alpha(t)$ s.t. $\alpha(0) = f_1$ and $\alpha(1) = f_2$. Let $q_i = Q(f_i), i = 1, 2$ be the SRNFs of the given surfaces f_1 and f_2 . Let $\beta : [0, 1] \rightarrow \mathbb{L}^2(\mathbf{S}^2, \mathbb{R}^3)$ denote the geodesic path, obtained via a straight line connecting q_1 and q_2 . Then, for any arbitrary point $\beta(\tau) \in \mathbb{L}^2(\mathbf{S}^2, \mathbb{R}^3)$, we want to find a surface $\alpha(\tau)$ such that




Fig. 2. Four examples of reconstructing anatomical surfaces. In each cell, the surface on the left is the ground truth (f_o) while the reconstruction (f^*) is on the right. The corresponding energies $\|Q(f^*) - Q(f_o)\|^2$ are shown at the top.




Fig. 3. Inversion from SRNF space to shape space gives an alternative way to analyze shapes. Previous methods require pulling back the metric and working with the Riemannian metric on \mathcal{F} (left). The proposed method performs analysis in \mathbb{L}^2 space (right) and pulls back the results onto the shape space (left) by inversion.

$Q(\alpha(\tau)) = \beta(\tau)$. In practice, we will accomplish this sequentially. For any $\epsilon > 0$, we start by solving for $f(\epsilon)$. Since our search is gradient-based, we need a good initial condition for starting the search. In this case $\alpha(0) = f_1$ provides such an initial condition. For the next shape, $f(2\epsilon)$, we can use the previous step $f(\epsilon)$ to initialize the search, and so on.

Figure 4 shows results of computing a geodesic connecting two known endpoints given by synthetic surfaces. The path of shapes is initialized by linear interpolation of SRNFs and then optimized numerically to form a geodesic path. An arbitrary path is shown to the right for comparison. Paths of energies are shown in the bottom panel. The energy paths of the arbitrary path, the linear path and the numerically computed geodesic path are shown in green, blue and red, respectively. We observe that the analytically inverted path has low energy and is close to the solution. The computed geodesic is shown in the left panel: it smoothly deforms one shape into the other. Similar experiments are performed with anatomical surfaces; the geodesics are shown in Fig. 5.




Fig. 4. Constructing geodesic paths connecting two shapes. The computed geodesic path is shown to the left compared to an arbitrary path to the right. In the bottom plot, the energy path along the geodesic is shown as the dash-dot line with circles (red), while that of the initialized linear path is the dashed line with triangles (blue). The energy along the arbitrary path is shown as a solid line with squares (green) as a comparison.




Fig. 5. Two geodesics computed for anatomical surfaces. Geodesics connecting the given two endpoint surfaces, f_1 and f_2 , are shown as $\alpha(t)$ at discrete time stamps for the thalamus and the pallidum.

2. **Shooting Geodesics:** Given a surface f and a tangent vector v_0 at f , one wants to construct a geodesic path $\alpha(t)$ s.t. $\alpha(0) = f$ and $\dot{\alpha}(0) = v$. Here $\dot{\alpha} = d\alpha/dt$. Note that shooting a geodesic is essentially evaluating the exponential map $\exp_f(tv_0) = \alpha(t), t = [0, 1]$ numerically. Let $\beta : [0, 1] \rightarrow \mathbb{L}^2(\mathbb{S}^2, \mathbb{R}^3)$ denote a straight line, i.e. $\beta(t) = Q(f) + tQ_{*,f}(v_0)$, where $Q_{*,f}$ is the differential of Q at f as previously mentioned. Then the desired geodesic path $\alpha(t)$ is of the form $Q(\alpha(t)) = \beta(t)$. This path $\alpha(t)$ is computed sequentially similarly to the first case. Some statistical analyses computed using shooting geodesics are shown in Fig. 6 and 7.
3. **Statistical Summaries of Shapes:** When given a sample of observed surfaces f_1, \dots, f_n , one wants to estimate the mean shape and principal directions of variation.

The mean shape μ_f is computed as shown in Table 1. Let $q_i, i = 1, \dots, n$, be the SRNFs of the registered surfaces in the sample and u_q^k be the k -th principal component of q_1, \dots, q_n . The k -th principal mode of variation for the SRNFs is

given by $\mu_q \pm \lambda u_q^k$, $\lambda \in \mathbb{R}^+$. In order to visualize the principal directions in the shape space, we need to find f^k such that $Q(f^k) = \mu_q \pm \lambda u_q^k$. This is essentially a shooting geodesic type of problem.

We generated two groups of synthetic surfaces, each with 8 observations, as shown in Fig. 6. Within each group, we computed the mean shape and performed principal component analysis. The first three principal directions (PD) are shown on the mean shapes of each group as their local magnitudes. Computed mean shapes and modes of variation on anatomical surfaces are presented in Fig. 7. Under the proposed framework, to compute the Karcher mean, computational cost per iteration (Algorithm 2 in Table 1) is 174 seconds comparing to 397 seconds in the previous method (Algorithm 1 in Table 1) using the PCA basis (8 of them, see [27]) and more than 4 hours using 200 spherical harmonic basis. Inverting the μ_q takes 6 seconds.

- 4. Random Sampling from Shape Models:** When given a sample of observed surfaces f_1, \dots, f_n , one wants to fit a probability model to the data and generate random samples from it. Let q_1, \dots, q_n be the SRNFs of the registered surfaces

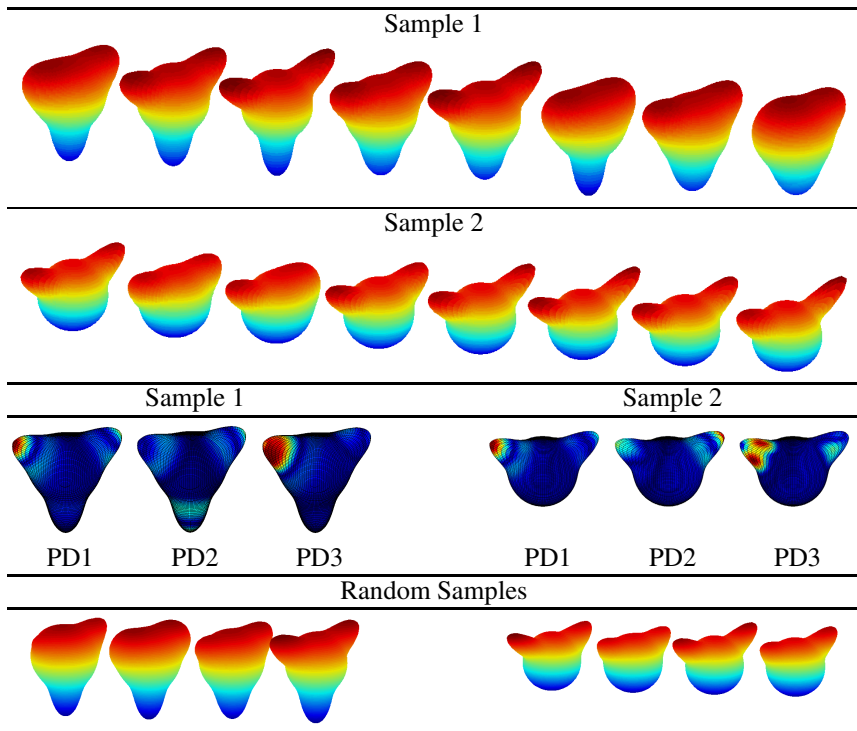


Fig. 6. Statistical analysis of synthetic data sets. Each sample has eight observations. The first three principal directions (PD) are shown plotted on the corresponding mean shapes for both samples (middle). Deformation magnitude is shown by color (blue small, red large). Random samples from Gaussian models are shown at the bottom.




Fig. 7. Plots of mean shape and principal directions (PD) for medical surfaces. Deformation magnitude is shown by color (blue small, red large) and plotted on mean shapes.




Fig. 8. Transfer of a deformation across shapes. Surfaces f_1 , h_1 and f_2 are given. Deformation from f_1 to h_1 is learnt and used to deform f_2 to get the new surface h_2 .

from the last step and $G(q)$ be the model probability distribution. A random sample can be generated from G and we denote it as q_s . We want to find f_s such that $Q(f_s) = q_s$ and it will be a randomly sampled shape. Using the registered SRNFs from Fig. 6, we used the principal components and estimated a multivariate Gaussian model for each group. Random samples of SRNFs are generated from the corresponding models and random shapes from both models are shown in the shape space by inversion in the bottom row of Fig. 6.

- Transferring Deformation between Shapes:** Given surfaces f_1 , h_1 and f_2 , one wants to find h_2 such that f_2 deforms to it in a similar way f_1 deforms to h_1 . In this case we are interested in estimating deformations between two shapes and then applying the deformations to new test shapes. The task can be decomposed into three components: (1) to learn the deformation from f_1 to h_1 as v , (2) to transfer v at f_1 to f_2 resulting v^{\parallel} and (3) to deform f_2 into h_2 using v^{\parallel} . Steps (1) and (3) are achieved by constructing geodesics while step (2) needs the tool of parallel transport. The detailed algorithm is described in Table 1. Figure 8 shows an example of transferring a deformation from one surface to another in the shape space.

5 ADHD Classification

In this section we apply our approach to an important problem in medical image analysis: the diagnosis of attention deficit hyperactivity disorder (ADHD) using MRI scans. The surfaces of brain structures used here were extracted from 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 [15,14,6]. Among the 34 subjects studied, 19 were diagnosed with ADHD and the remaining 15 were controls (non-ADHD). Some examples of left structures are displayed in Fig. 9. First




Fig. 9. Left anatomical structures in the brain

Table 2. Classification Performance for Five Different Techniques

Method Structure (%)	SRNF Gauss Proposed	SRM Gauss [19]	SRM NN [18]	Harmonic	ICP [3]	SPHARM PDM [24]
L. Caudate	67.7	-	41.2	64.7	32.4	61.8
L. Pallidus	85.3	88.2	76.5	79.4	67.7	44.1
L. Putamen	94.1	82.4	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

we register the extracted surfaces as described in [16] and map them into the \mathbb{L}^2 space of SRNFs using Q . In order to distinguish ADHD and control samples, we use the Gaussian classifier on principal components as defined in Section 4.

Table 2 shows the single structure, LOO nearest neighbor classification rate in %. The best performance is attained using the proposed SRNF Gaussian classifier between left putamen surfaces. We compare our results to those obtained using: the SRM Gaussian classifier; the SRM NN classifier; the iterative closest point (ICP) algorithm; an approach using fixed surface parametrization and \mathbb{L}^2 distance between the surfaces; and the SPHARM-PDM approach. The performance measures for these approaches were taken from Kurtek et al. [19] and other previously published papers. The results suggest that the parametrization-invariant metric and the probability models in our approach provides improved matching and modeling of the surfaces, resulting in a superior ADHD classification. In summary, our method is not only more efficient: the computational cost is an order of magnitude less than SRM and related ideas; but also provides significantly improved classification.

6 Conclusions

The SRNF representation is potentially an important tool in statistical shape analysis of parameterized surfaces. Previous methods built tools for analysis directly in the surface space, which is computationally inefficient. We have introduced methods for approximating the inverse mapping Q^{-1} . This map can be used to convert results computed in SRNF space back to the shape space. Since the SRNF space is a vector space with \mathbb{L}^2 -metric, the cost of statistical analysis in this space is very low, thus simplifying typical shape analysis tasks. In general, by adopting the proposed framework, computational

cost of algorithms for various analyses can be reduced by an order of magnitude. Experimental results show that the same analyses can be performed under the simplified framework, and that the method achieves state-of-the-art performance on the classification of ADHD data.

Acknowledgement. This research was supported in part by the NSF grants DMS 1208959, IIS 1217515, and CCF 1319658. We also thank the producers of datasets used here for making them available to public.

References

1. Abe, K., Erbacher, J.: Isometric immersions with the same Gauss map. *Mathematische Annalen* 215(3), 197–201 (1975)
2. Allen, B., Curless, B., Popović, Z.: The space of human body shapes: Reconstruction and parameterization from range scans. *ACM Transactions on Graphics* 22(3), 587–594 (2003)
3. Besl, P.J., McKay, N.D.: A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2), 239–256 (1992)
4. Bouix, S., Pruessner, J.C., Collins, D.L., Siddiqi, K.: Hippocampal shape analysis using medial surfaces. In: Niessen, W.J., Viergever, M.A. (eds.) *MICCAI 2001*. LNCS, vol. 2208, pp. 33–40. Springer, Heidelberg (2001)
5. Brechbühler, C., Gerig, G., Kübler, O.: Parameterization of closed surfaces for 3D shape description. *Computer Vision and Image Understanding* 61(2), 154–170 (1995)
6. Burden, M.J., Jacobson, J.L., Westerlund, A., Lundahl, L.H., Morrison, A., Dodge, N.C., Klorman, R., Nelson, C.A., Avison, M.J., Jacobson, S.W.: An event-related potential study of response inhibition in ADHD with and without prenatal alcohol exposure. *Alcoholism: Clinical and Experimental Research* 34(4), 617–627 (2010)
7. Davies, R.H., Twining, C.J., Cootes, T.F., Taylor, C.J.: Building 3-D statistical shape models by direct optimization. *IEEE Transactions on Medical Imaging* 29(4), 961–981 (2010)
8. Dryden, I., Mardia, K.: *Statistical Shape Analysis*. John Wiley & Sons (1998)
9. Eschenburg, J.-H., Kruglikov, B.S., Matveev, V.S., Tribuzy, R.: Compatibility of Gauss maps with metrics. *Differential Geometry and its Applications* 28(2), 228–235 (2010)
10. Gorczowski, K., Styner, M., Jeong, J.Y., Marron, J.S., Piven, J., Hazlett, H.C., Pizer, S.M., Gerig, G.: Multi-object analysis of volume, pose, and shape using statistical discrimination. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32(4), 652–661 (2010)
11. Grenander, U., Miller, M.I.: Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics* LVI(4), 617–694 (1998)
12. Heeren, B., Rumpf, M., Wardetzky, M., Wirth, B.: Time-discrete geodesics in the space of shells. *Computer Graphics Forum* 31(5), 1755–1764 (2012)
13. Hoffman, D.A., Osserman, R.: The Gauss map of surfaces in R3 and R4. *Proceedings of the London Mathematical Society* 3(1), 27–56 (1985)
14. Jacobson, S.W., Jacobson, J.L., Sokol, R.J., Chiodo, L.M., Corobana, R.: Maternal age, alcohol abuse history, and quality of parenting as moderators of the effects of prenatal alcohol exposure on 7.5-year intellectual function. *Alcoholism: Clinical and Experimental Research* 28(11), 1732–1745 (2004)
15. Jacobson, S.W., Jacobson, J.L., Sokol, R.J., Martier, S.S., Chiodo, L.M.: New evidence for neurobehavioral effects of in utero cocaine exposure. *The Journal of Pediatrics* 129(4), 581–590 (1996)

16. Jermyn, I., Kurtek, S., Klassen, E., Srivastava, A.: Elastic shape matching of parameterized surfaces using square root normal fields. In: IEEE European Conference on Computer Vision, vol. 5(14), pp. 805–817 (2012)
17. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: ACM SIG-GRAPH (2007)
18. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J., Avison, M., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. IEEE Transactions on Medical Imaging 30(3), 849–858 (2011)
19. Kurtek, S., Klassen, E., Ding, Z., Avison, M.J., 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)
20. Kurtek, S., Klassen, E., Gore, J.C., Ding, Z., Srivastava, A.: Elastic geodesic paths in shape space of parameterized surfaces. IEEE Transactions on Pattern Analysis and Machine Intelligence 34(9), 1717–1730 (2012)
21. Osher, S., Fedkiw, R.: Level Set Methods and Dynamic Implicit Surfaces (Applied Mathematical Sciences). 2003 edn. Springer(November 2002)
22. Ovsjanikov, M., Li, W., Guibas, L., Mitra, N.J.: Exploration of continuous variability in collections of 3D shapes. ACM Transactions on Graphics 30(4), 33:1–33:10 (2011)
23. Srivastava, A., Klassen, E., Joshi, S.H., Jermyn, I.H.: Shape analysis of elastic curves in Euclidean spaces. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(7), 1415–1428 (2011)
24. Styner, M., Oguz, I., Xu, S., Brechbuehler, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. Insight Journal (July 2006)
25. Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A Survey on Shape Correspondence. Computer Graphics Forum 30(6), 1681–1707 (2011)
26. Windheuser, T., Schlickewei, U., Schmidt, F.R., Cremers, D.: Geometrically consistent elastic matching of 3D shapes: A linear programming solution. In: IEEE International Conference on Computer Vision, pp. 2134–2141 (November 2011)
27. Xie, Q., Kurtek, S., Le, H., Srivastava, A.: Parallel transport of deformations in shape space of elastic surfaces. In: IEEE International Conference on Computer Vision (December 2013)
28. Yang, Y.L., Yang, Y.J., Pottmann, H., Mitra, N.J.: Shape space exploration of constrained meshes. ACM Transactions on Graphics 30(6), 124:1–124:12 (2011)
29. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., Van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. Computer Graphics Forum 27(5), 1431–1439 (2008)