

Numerical Inversion of SRNF Maps for Elastic Shape Analysis of Genus-Zero Surfaces

Hamid Laga, Qian Xie, Ian H. Jermyn *Member, IEEE*, and Anuj Srivastava *Fellow, IEEE*

Abstract—Recent developments in elastic shape analysis (ESA) are motivated by the fact that it provides a comprehensive framework for simultaneous registration, deformation, and comparison of shapes. These methods achieve computational efficiency using certain square-root representations that transform invariant elastic metrics into Euclidean metrics, allowing for the application of standard algorithms and statistical tools. For analyzing shapes of embeddings of S^2 in \mathbb{R}^3 , Jermyn et al. [1] introduced square-root normal fields (SRNFs), which transform an elastic metric, with desirable invariant properties, into the \mathbb{L}^2 metric. These SRNFs are essentially surface normals scaled by square-roots of infinitesimal area elements. A critical need in shape analysis is a method for inverting solutions (deformations, averages, modes of variations, etc.) computed in SRNF space, back to the original surface space for visualizations and inferences. Due to the lack of theory for understanding SRNF maps and their inverses, we take a numerical approach, and derive an efficient multiresolution algorithm, based on solving an optimization problem in the surface space, that estimates surfaces corresponding to given SRNFs. This solution is found to be effective even for complex shapes that undergo significant deformations including bending and stretching, *e.g.* human bodies and animals. We use this inversion for computing elastic shape deformations, transferring deformations, summarizing shapes, and for finding modes of variability in a given collection, while simultaneously registering the surfaces. We demonstrate the proposed algorithms using a statistical analysis of human body shapes, classification of generic surfaces, and analysis of brain structures.

Index Terms—Elastic shape analysis, square-root representations, Riemannian metrics, elastic registration, shape statistics, shape modeling.



1 INTRODUCTION

Shape analysis is an important area of research with a wide variety of applications. The need for shape analysis arises in many branches of science, including anatomy, bioinformatics, computer graphics, and 3D printing and prototyping. A comprehensive method for shape analysis provides: (i) a **shape metric**, for pairwise quantification of shape differences between objects; (ii) a **shape summary**, a compact representation of a class of objects in terms of the center (mean or median) and covariance of the shapes; and (iii) **modes of variability**, the dominant modes of variation around the center, analogous to PCA, but for modeling shape variability.

Although shape analysis is a broad area that potentially involves comparisons of objects with different topologies as well as different geometries, we will restrict ourselves to geometric variability by choosing a fixed topology; specifically, we focus on shapes of genus-0 surfaces. This is still a very rich class of shapes with tremendous variability, and applications

in many scientific disciplines. Hereafter, we will highlight the growing field of elastic shape analysis of genus-0 surfaces, and isolate an important problem that becomes the focus of this paper.

1.1 Previous Methods

Due to the importance of shape analysis, there is a large body of previous work addressing it. However, most of this work shares a common drawback, as described next. Generally speaking, there are two subproblems of interest when comparing two shapes: **registration** and **deformation**. Registration deals with complete or partial matching of points across objects, *i.e.* deciding which point on one object matches which point on the other; while deformation is concerned with finding a continuous sequence of shapes starting from one shape and ending at the other. Most existing techniques treat registration and deformation as independent problems, despite their obvious interdependence; for example, the optimality of the solutions to each of the subproblems may be governed by different objective functions or metrics. We now discuss some examples of previous work, pointing out where this occurs.

A prominent statistical shape analysis framework, pioneered by Kendall's school [2], [3], works with point sets that are already registered, and focuses only on deformations. Approaches such as medial surfaces [4], [5] and level sets [6] either presume registration, or solve for it using some independent pre-processing criterion such as MDL [7]. Kilian *et al.* [8]

- Hamid Laga is with the School of Engineering and IT, Murdoch University (Australia), and with the Phenomics and Bioinformatics Research Centre, University of South Australia. E-mail: H.Laga@murdoch.edu.au
- Qian Xie is with the Department of Statistics, Florida State University.
- Ian H. Jermyn is with the Department of Mathematical Sciences, Durham University (UK).
- Anuj Srivastava is with the Department of Statistics, Florida State University.

represent surfaces by discrete triangulated meshes and compute geodesic paths (deformations) between them, while assuming that the meshes are registered. Heeren *et al.* [9] propose a method for computing geodesic-based deformations of thin shell shapes, with extensions for computing summary statistics in the shell space [10], but with known registration.

Some other papers solve for registration [11], [12] while ignoring deformation. Examples include Windheuser *et al.* [13], who solves a dense registration problem, but uses linear interpolation between registered points in \mathbb{R}^3 to form deformations. Techniques such as SPHARM or SPHARM-PDM [14], [15] seek uniform sampling on the domain to address the registration issue. This is a major restriction, as it limits registration of corresponding features across surfaces.

In conclusion, a majority of the papers in the literature treat the registration and deformation problems in a disjoint fashion, each with their own optimality criteria. Due to this disconnection, the overall shape analysis pipeline becomes suboptimal. As an exception, Hirshberg *et al.* [?] present an interesting method for analyzing human body shapes using an energy-minimization approach for joint registration and modeling of deformation. However, since this energy function is not a formal metric, it is not possible to derive a compatible statistical analysis that results in means, covariances, principal components, or parallel transports of deformations.

1.2 Elastic Shape Analysis and SRNFs

A more comprehensive approach to shape analysis is to perform registration and deformation in a single **unified** and **metric-based** fashion. Elastic shape analysis (ESA) is a Riemannian approach that accomplishes exactly that, for objects such as curves and surfaces. In ESA, one identifies an appropriate representation space for parameterized versions of the objects, and endows it with a Riemannian metric.¹ The metric allows the definition of a *geodesic*, *i.e.* a shortest path, between two objects, which is by definition the optimal deformation under the metric. Registration is handled by finding the shortest path, not between the original two parameterized objects, but between all possible reparameterizations of those objects. This simultaneously identifies the optimal registration of the objects, and renders the resulting registration and deformation independent of the original parameterizations. The result is therefore a registration and deformation of the corresponding *geometric* (*i.e.* unparameterized) objects. The optimization over all possible reparameterizations can alternatively be viewed as finding a geodesic in a quotient of the original space by the group of reparameterizations, using a

Riemannian metric derived from that on the original space. One can also choose to introduce further invariances to, for example, translations and rotations, thereby incorporating alignment into the deformation.

Central to this program is the definition of a Riemannian metric on the space of parameterized objects that is preserved by the action of the relevant transformations, typically reparameterizations, translations, rotations, and perhaps scale. Key to its success in practical terms is the definition of a Riemannian metric that ‘measures’ the types of shape changes that are important, and that renders feasible computationally the corresponding geodesic calculations.

These goals have been achieved in the shape analysis of **curves** by using a particular member of the family of *elastic metrics*, in conjunction with a representation called the *square-root velocity function* (SRVF). The form of the elastic metric is extremely simple when expressed in terms of the SRVF: it becomes the \mathbb{L}^2 metric [16]. This greatly facilitates computations, enabling sophisticated statistical analyses that require many geodesic calculations; consequently, this method has been used in many practical problems. 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 be uniquely mapped back to the original curve space, using an analytical expression. This is significantly more efficient than performing analysis in the curve space itself.

To analyze genus-0 **surfaces**, Kurtek *et al.* [17] introduced the *square-root map* (SRM). 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 $s \in \mathbf{S}^2$, then the SRM of f is defined to be $q(s) = |n(s)|^{1/2} f(s)$, where $n(s) = f_u(s) \times f_v(s)$ is the unnormalized normal to the surface at $f(s)$. Here $f_u = \partial f / \partial u$ and $f_v = \partial f / \partial v$. The metric is then taken to be the \mathbb{L}^2 metric in SRM space. Unfortunately, this metric has several limitations, including that the metric distance between two shapes changes when they are both translated by the same amount.

Jermyn *et al.* [1] addressed these issues by defining a new representation termed the *square-root normal field* (SRNF). A parameterized surface $f \mapsto Q(f)$, where $Q(f)(s) \equiv n(s)/|n(s)|^{1/2}$. The authors showed that the Euclidean distance between two points in the space of SRNFs is equivalent to geodesic distance under an elastic metric on surfaces [1], *i.e.* the metric is interpretable as the weighted sum of the bending and stretching needed to deform the corresponding surfaces into one another. Thus, as in the case of the SRVF for curves, equipping the space of SRNFs, Q , with the \mathbb{L}^2 metric means that many computations, *e.g.* those for the mean, covariance, and PCA, are greatly simplified.

We add some details to explain these ideas further.

1. We refer the reader to http://stat.fsu.edu/~anuj/CVPR_Tutorial/ for a comprehensive tutorial on differential geometry methods in shape analysis.

Let Γ be the group of all orientation-preserving diffeomorphisms of \mathbf{S}^2 , and $SO(3)$ be the group of all rotations in \mathbb{R}^3 . The product group $\mathcal{G} \equiv SO(3) \times \Gamma$ helps align and register surfaces as follows. For any two parameterized surfaces $f_1, f_2 \in \mathcal{F}$, and $s \in \mathbf{S}^2$, the points $f_1(s)$ and $f_2(s)$ are considered paired or matched originally. However, for any $(O, \gamma) \in \mathcal{G}$, the surface $O(f_2 \circ \gamma)$ denotes a rotated and reparameterized version of f_2 , and now $f_1(s)$ is paired with $O(f_2(\gamma(s)))$. Thus, the elements of \mathcal{G} can be used to control the (rotational) alignment and registration of surfaces. The SRNF of the transformed surface is given by $O(Q(f), \gamma)$ where $(q, \gamma)(s) \equiv q(\gamma(s))\sqrt{J_\gamma(s)}$ and $J_\gamma(s)$ denotes the determinant of the Jacobian of γ at s . This is formally an action of \mathcal{G} on the set \mathbb{L}^2 . The most important property of SRNFs is that, for any $f_1, f_2 \in \mathcal{F}$ and $(O, \gamma) \in \mathcal{G}$, we have:

$$\|O(Q(f_1), \gamma) - O(Q(f_2), \gamma)\| = \|Q(f_1) - Q(f_2)\|, \quad (1)$$

where $\|\cdot\|$ denotes the \mathbb{L}^2 norm. In other words, the group \mathcal{G} acts by isometries under the \mathbb{L}^2 metric in \mathcal{Q} , as required. Note that the SRNF representation is translation invariant by definition because defined in terms of derivatives. As stated above, the \mathbb{L}^2 metric on \mathcal{Q} corresponds to a partial elastic Riemannian metric on \mathcal{F} [1].

The space \mathcal{Q} is often called *pre-shape* space, since many elements correspond to the same geometric object. For instance, $Q(f)$ and $O(Q(f), \gamma)$ can be two different elements of \mathcal{Q} , but represent rotated versions of the same surface. To facilitate ESA, these representations are identified using an equivalence relation, each shape being represented by an orbit under \mathcal{G} :

$$[Q(f)] = \{O(Q(f), \gamma) | (O, \gamma) \in \mathcal{G}\}. \quad (2)$$

The set of all such orbits is the quotient space $\mathcal{S} = \mathcal{Q}(\mathcal{F})/\mathcal{G}$, also called *shape space*. The computation of geodesics in \mathcal{S} thereby corresponds to joint registration and deformation of surfaces.

1.3 Open Issues and Problem Motivation

The SRNF is very promising as a representation of surfaces for ESA. If 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, there would be large gains in computational efficiency with respect to *e.g.* [1], [10], [18]. Unfortunately, the SRNF shares one difficulty with the SRM, and that is the problem of inversion. First, unlike the curve case, there is no analytical expression for Q^{-1} for arbitrary points in \mathbb{L}^2 . Moreover, the injectivity and surjectivity of Q remain to be determined, meaning that for a given $q \in \mathcal{Q}$, there may be no $f \in \mathcal{F}$ such that $Q(f) = q$, and if such an f does exist, it may not be unique.² If one cannot

2. Note that 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.

invert the representation, it is not clear how to transfer geodesics and statistical analyses conducted in \mathcal{Q} back to \mathcal{F} , thereby removing one of the main motivations for introducing the SRNF. One can always pull the \mathbb{L}^2 metric back to \mathcal{F} under Q and perform computations there, as in [18], but this is computationally expensive, and rather defeats the purpose of having an \mathbb{L}^2 metric in the first place.

In this paper, we address this problem by developing a method that, given $q \in \mathcal{Q}$, finds an $f \in \mathcal{F}$ such that $Q(f) = q$, if one exists, or an f whose image $Q(f)$ is the closest, in the elastic metric, to q if it does not. We achieve this by formulating SRNF inversion as an optimization problem: find an element $f \in \mathcal{F}$ whose image $Q(f)$ is as close as possible to the given $q \in \mathcal{Q}$ under the \mathbb{L}^2 norm. We then propose an efficient numerical procedure to solve this problem. In particular, we show that by carefully engineering an orthonormal basis of \mathcal{F} , and combining it with a spherical-wavelet based multiresolution and multiscale representation of the elements of \mathcal{F} , SRNF maps can be inverted with high accuracy and efficiency, with robustness to local minima, and with low computational complexity. We also show that an *analytical* solution to the inversion problem does nevertheless exist for star-shaped surfaces, *i.e.* those whose enclosed volumes are star domains, a large family of surfaces found in many real problems.

Using the proposed methods for inverting the SRNF map, we are able to compute geodesics, transfer deformations, perform statistical analyses, and synthesize random shapes sampled from the space of genus-0 surfaces, in a very efficient manner: the **computational cost is reduced by an order of magnitude** compared to previous work. We demonstrate the utility of the proposed tools using complex genus-0 surfaces exhibiting complex bending and stretching deformations, taken from the TOSCA dataset, the SHREC07 watertight database [19], and the human shape database of [20].

The paper is organized as follows. Section 2 describes the inversion problem, its formulation as an optimization problem, and its numerical solution, as well as the analytic solution in the case of star-shaped surfaces. To demonstrate the ESA framework, Section 3 lists some important statistical tasks, and compares the algorithms for these tasks under previous and the proposed methods. Section 4 presents experimental results on computing geodesics; transferring deformations; computing statistical summaries; and performing classification of generic 3D shapes, and of subjects with Attention Deficit Hyperactivity Disorder using the shapes of brain subcortical structures.

2 THE INVERSION PROBLEM

In order to take full advantage of the simplicity of the metric in the SRNF representation, given $q \in \mathcal{Q}$, we

need to be able to find a surface f , if it exists, such that $Q(f) = q$. This is the inversion problem. Let $g(s)$ denote the 2×2 matrix $\begin{bmatrix} \langle f_u, f_u \rangle & \langle f_u, f_v \rangle \\ \langle f_u, f_v \rangle & \langle f_v, f_v \rangle \end{bmatrix}$ at each $s \in \mathbf{S}^2$. The inversion problem can be formulated in several different ways:

- 1) Given a map $\tilde{n} : \mathbf{S}^2 \rightarrow \mathbf{S}^2$ and a function $\omega : \mathbf{S}^2 \rightarrow \mathbb{R}_+$, when are these the Gauss map $\tilde{n} = n/|n|$ and area form $\omega = |n|^{1/2} = \det(g(s))^{1/2}$ of an immersion $f : \mathbf{S}^2 \rightarrow \mathbb{R}^3$? If such an f exists, is it unique (up to translation)?
- 2) Given $q : \mathbf{S}^2 \rightarrow \mathbb{R}^3$, does the differential equation $n = |q|q$ have a (global) solution, and is it unique (up to translation) [21]?

In the first case, it is known that if we are given (\tilde{n}, g) , rather than just (\tilde{n}, ω) , then f , if it exists, can be reconstructed up to translation and rotation [22], [23]. This is also true when we are given the conformal class of g [24], *i.e.* we know (\tilde{n}, \tilde{g}) , where $\tilde{g}(s) = g(s)/\det(g(s))^{1/2}$. Unfortunately, little seems to be known about the nature of the solution in the case at hand. The problem was stated by Arnold in 1990 [25], but in a 2004 book listing Arnold's problems and their subsequent elucidation or solution [26], the same problem still appears, without commentary.

In the second case, a global solution clearly does not always exist, at least in degenerate cases (if g is constant, for example [21]). The discussion in [21] does lead to the conclusion that the equation is locally solvable, but this is not sufficient for us: we require a global solution. Even if a global solution does exist, it may not be unique. In degenerate cases, any flat area of the surface could be subjected to an area-preserving diffeomorphism, while for open surfaces, one may find distinct, non-degenerate cases that share the same Gauss map and area form [27]. If the solution is to be unique, then it can be so only for non-degenerate, closed surfaces.

We will manage these difficult questions about the injectivity and surjectivity of Q by formulating inversion as an optimization problem. We define an energy function $E_0 : \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$ by

$$E_0(f; q) = \inf_{(O, \gamma) \in \mathcal{G}} \|Q(f) - O(q, \gamma)\|_2^2, \quad (3)$$

where $(q, \gamma) = (q \circ \gamma)\sqrt{J_\gamma}$, and J_γ is the determinant of the Jacobian of γ . Finding an $f \in \mathcal{F}$ such that $Q(f) = q$ is then equivalent to seeking zeros of E_0 . We denote by f^* an element of \mathcal{F} satisfying $E_0(f^*; q) = \inf_{f \in \mathcal{F}} E_0(f; q)$.

The minimization will be performed using a gradient descent approach. From a computational point of view, it will be easier to deal with deformations of a surface, rather than the surface itself. We therefore set $f = f_0 + w$, where f_0 denotes the current estimate of f^* , and w is a deformation of f_0 . Then, we minimize

$$E(w; q) = \inf_{(O, \gamma) \in \mathcal{G}} \|Q(f_0 + w) - O(q, \gamma)\|_2^2, \quad (4)$$

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 $f_0 = 0$.

Formulating the inversion problem as energy minimization has two purposes. First, by starting gradient descent for similar $q, q' \in \mathcal{Q}$ from similar initial points $f_0, f'_0 \in \mathcal{F}$, we can help to ensure that the inverted surfaces we find are, if not uniquely determined, at least consistent with each other. This will be important when we wish to invert, for example, whole geodesics. Second, because there is no guarantee that Q is surjective, and, equally, no guarantee that a geodesic between two points in the image of Q lies wholly in the image of Q , we have to be able to 'invert' points in \mathcal{Q} that do not lie in the image of Q . Energy minimization achieves this by finding an f^* whose image under Q is as close as possible to the point whose inverse is desired. This process is equivalent to projecting points in \mathcal{Q} to $Q(\mathcal{F})$.

2.1 Numerical Optimization

We solve the optimization problem of Eqn. (4) by iterating over the following two steps:

- 1) Given f_0 and w , find a rotation O and a diffeomorphism γ such that the energy E is minimized. Note this step is not strictly necessary, since \mathcal{F} is closed under the action of these groups. Separating out these transformations, however, enables large 'leaps' in \mathcal{F} , helping the algorithm to search more globally.
- 2) Given a fixed O and γ , optimize E with respect to w using gradient descent.

The first step is solved using the approach described in [1]. That is, we find the best rotation O using SVD, given a fixed γ . Then we find the best γ by a search over the space of all diffeomorphisms using a gradient descent approach. In practice, we have found that restricting the search over Γ to the space of rigid re-parameterizations, $\Gamma_0 \equiv SO(3)$, is sufficient to achieve good reconstruction accuracy, and is also computationally very efficient since it can be solved using the Procrustes method. We emphasize that we only do this for the inversion process; we still use the full Γ for the registration step when computing geodesics between shapes.

To minimize E with respect to w using a gradient descent approach, we need the derivative of E . Since \mathcal{F} is an infinite-dimensional vector space, we will approximate the derivative using a finite basis for \mathcal{F} . We therefore set $w = \sum_{b \in \mathcal{B}} c_b b$, with $c_b \in \mathbb{R}$, and where \mathcal{B} forms an orthonormal basis of \mathcal{F} . The directional derivative of E at $f_0 + w$ in the direction

of b , $c_b \equiv \nabla_b E(w; q)$, is then given by:

$$\begin{aligned} \nabla_b E(w; q, f_0) &= \frac{d}{d\epsilon} \Big|_{\epsilon=0} \|Q(f_0 + w + \epsilon b) - q\|_2^2 \\ &= 2\langle Q(f_0 + w) - q, Q_{*,f_0+w}(b) \rangle. \end{aligned} \quad (5)$$

(Note that here, for notational simplicity, we have dropped O and γ .) Here $Q_{*,f}$ denotes the differential of Q at f . This can be evaluated using the following expression: for all $s \in \mathbf{S}^2$,

$$Q_{*,f}(b)(s) = \frac{n_b(s)}{\sqrt{|n(s)|}} - \frac{n(s) \cdot n_b(s)}{2|n(s)|^{5/2}} n(s), \quad (6)$$

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

$$Q_{*,f}(b)(s) = \frac{1}{\sqrt{|n(s)|}} \left(n_b(s) - \frac{\tilde{n}(s) \cdot n_b(s)}{2} \tilde{n}(s) \right). \quad (7)$$

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. 5, 6, and 7.

A naive implementation of this procedure results in the algorithm becoming trapped in local minima. We show, however, that this problem can be overcome by carefully engineering the orthonormal basis \mathcal{B} of \mathcal{F} (Section 2.1.1) and combining it with a multiscale and multiresolution representation of the elements of \mathcal{F} and \mathcal{Q} (Sections 2.1.2 and 2.1.3).

2.1.1 Basis Formation

An important ingredient of our approach is the representation of the vector space \mathcal{F} using an orthonormal basis set \mathcal{B} . Since elements of \mathcal{F} are smooth mappings of the type $f : \mathbf{S}^2 \rightarrow \mathbb{R}^3$, we can use spherical harmonics (SHs), in each coordinate, to form the basis set. Let $\{Y_i(s)\}$, for i in some index set, be the real-valued SH functions. Then, a basis for representing surfaces $f : \mathbf{S}^2 \rightarrow \mathbb{R}^3$ can be constructed as $\{Y_i e_1\} \cup \{Y_j e_2\} \cup \{Y_k e_3\}$, where e_1, e_2, e_3 is the standard basis for \mathbb{R}^3 . We will refer to the resulting basis as $\mathcal{B} = \{b_j\}$, for j in some index set. \mathcal{B} is a generic basis that can be used to represent arbitrary spherically-parameterized surfaces. However, complex surfaces will require a large number of basis elements to achieve high accuracy and capture all the surface details. We show that by using a spherical wavelet-based multiresolution approach, the optimization problem of Eqn. (3) can be solved in a computationally efficient manner.

Alternatively, depending on the context and the availability of training data, it may be more efficient to use a principal component basis instead of a generic basis. For instance, in the case of human shapes, there are several publicly available data sets that can be used to build a PCA basis. Thus, given a set of training samples, we can compute its PCA in \mathcal{F} and use the first N components to construct the basis. In our experiments, the number of PCA basis elements

required is of the order of 100, significantly smaller than the number of SH elements (more than 3000).

Finally, note that one can also combine the PCA and SH bases in order to account for shape deformations that have not been observed in the training dataset. It is also possible to use different bases for the surfaces (elements of \mathcal{F}) and the deformations (elements of tangent spaces $T_f \mathcal{F}$) to improve efficiency, although we do not explore that here.

2.1.2 Multiscale representation of SRNFs

The gradient descent method at a fixed resolution finds it difficult to reconstruct complex surfaces that contain elongated parts and highly non-convex regions, such as the ones shown in Fig. 2. We observe however that such parts often correspond to the high frequency components of the surface, which can be easily captured using a multiscale analysis.

Given an SRNF q of an unknown surface f , and an SRNF q_s of a unit sphere f_s , the geodesic between them under the \mathbb{L}^2 metric is the linear path $\beta(\tau) = (1 - \tau)q_s + \tau q, \tau \in [0, 1]$; let the corresponding path in \mathcal{F} be $\alpha(\tau)$ such that $f_s = \alpha(0)$ and $f = \alpha(1)$. One can then treat β as a multiscale representation of the SRNF q . In practice, we divide the linear path β into $m + 1$ equidistant points $q^i, i = 0 \dots, m$, such that $q^0 = q_s, q^m = q$, and $\|q^i - q^{(i-1)}\| = \Delta > 0$. We refer to $\mathbf{q} = (q^1, \dots, q^m)$ as the multiscale SRNF of an unknown multiscale surface $\mathbf{f} = (f^1, \dots, f^m)$. (Here, $f^0 = f_s$ and $q^0 = q_s$ are not included in the representation.)

With this representation, the SRNF inversion problem can be reformulated as follows. Given a multiscale SRNF $\mathbf{q} = (q^1, \dots, q^m)$, the goal is to find a multiscale surface $\mathbf{f} = (f^1, \dots, f^m)$ such that $Q(f^i)$ is as close as possible to q^i , i.e. $E_0(f^i; q^i)$ is minimized. Gradient descent procedures are known to provide good results if the initialization is close to the solution. We thus use the following iterative procedure: in step 1, we find a surface \tilde{f}^1 that minimizes $E_0(f; q^1)$, initializing gradient descent with a sphere f^0 . Then, at step i , we find the surface \tilde{f}^i that minimizes $E_0(f; q^i)$, initializing gradient descent with $\tilde{f}^{(i-1)}$. Fig. 3 of the supplementary material shows a few examples of SRNF inversion using this multiscale approach.

2.1.3 Multiresolution Inversion Algorithm

While this enables us to invert a single SRNF, using the gradient descent-based optimization procedure to converge to the desired solution, it can be computationally very expensive when dealing with high resolution surfaces. In practice, the computation time can be significantly reduced by adopting a multiresolution representation of the elements of \mathcal{F} , and subsequently the elements of \mathcal{Q} . Specifically, given a surface $f : \mathbf{S}^2 \rightarrow \mathbb{R}^3$, we use spherical wavelet analysis to build a multiresolution representation $\mathbf{f} = (f^1, \dots, f^m = f)$ of f . The surface f^i at level i is a smoothed, sub-sampled, and thus coarse version of the surface f^{i+1} .

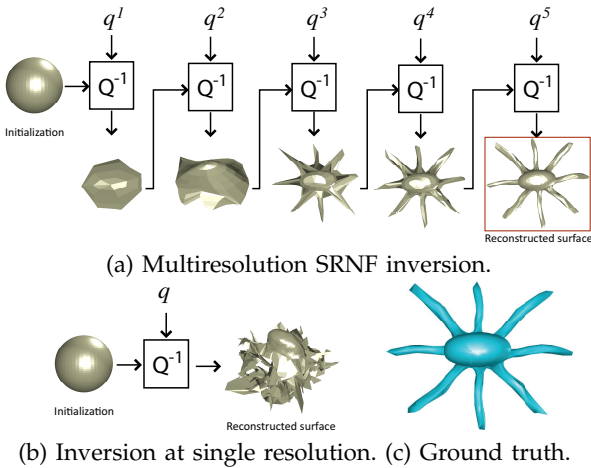


Fig. 1: Comparison between the proposed multiresolution approach (a) with the single resolution version (b). The latter fails to recover complex surfaces.

With a slight abuse of notation, we also define the multiresolution SRNF map sending $\mathbf{f} \mapsto Q(\mathbf{f}) = \mathbf{q} = (q^i)$, where $q^i = Q(f^i)$.

To build the multiresolution representation, we apply an analysis filter \tilde{h}_i to the spherical surface f via spherical convolution, resulting in the output $w_i = f \star \tilde{h}_i$. We then convolve w_i with another spherical filter h_i , a synthesis filter, producing a coarse surface f^i that is an approximation of the surface f at scale i . By defining the analysis filters $\{\tilde{h}_i\}$ as dilated versions of a template filter ψ , *i.e.* $\tilde{h}_i = D_i\psi$, the coefficients $\{w_n\}$ and thus the reconstructed surfaces $\{f^i\}$ capture the original surface properties at multiple scales. For details of these filters, we refer the reader to [28]. In practice, we choose the template wavelet ψ to be the Laplacian-of-Gaussian and we compute the wavelet decomposition of each coordinate function using convolutions with the analysis-synthesis filters in the SH domain. Since surfaces at large scales require a smaller number of samples, we start with a spherical grid of resolution 128×128 and at each wavelet decomposition we subsample the obtained coarse surface by halving the spherical resolution. We have found that four decomposition levels are sufficient to achieve high accuracy inversion.

With this representation, a multiresolution SRNF can be inverted using the same procedure as the one described in Section 2.1.2. This time, the surface obtained at each iteration is first up-sampled and then used to initialize optimization at the next resolution. Fig. 1(a) shows a reconstruction of the octopus surface from its SRNF using the proposed multiresolution inversion procedure. Fig. 1(b) displays the inversion result when Eqn. 4 is solved using a single-resolution representation. The latter results in a degenerate surface since the initialization is too far from the correct solution. The proposed multiresolution procedure converges to the correct solution since at each opti-

mization step, the initialization is close to the correct solution. Note also that the multiresolution procedure significantly reduces computational cost: although the optimizations at coarser scales require a large number of iterations, they operate on spherical grids of low resolution (*e.g.* 16×16) and are thus very fast. At the finest scale, although the resolution is high (128×128), gradient descent only requires a few iterations since the initialization is already very close to the solution.

We point out that our approach assumes that \mathbf{q} , the multiresolution representation of the SRNF q of an unknown surface f , is available. This is not a restriction, since in practice, we are interested in computing geodesics between known surfaces f_1 and f_2 , deforming a given surface f_1 along a given geodesic direction, or computing summary statistics of a given set of surfaces. In the first application, for example, one can: (1) build a multiresolution representation of the surfaces f_1 and f_2 , hereinafter denoted by \mathbf{f}_1 and \mathbf{f}_2 , respectively, using the spherical wavelet decomposition described in this section; (2) compute their multiresolution SRNFs \mathbf{q}_1 and \mathbf{q}_2 ; and finally (3) invert the multiresolution SRNF $\mathbf{q} = (q^1, \dots, q^m)$ where $q^i = (1 - \tau)q_1^i + \tau q_2^i$, $\tau \in [0, 1]$. Section 3 elaborates this point further in different contexts.

2.2 Star-Shaped Surfaces

Here we consider a special subset of genus-0 surfaces, *star-shaped* surfaces, which are of great relevance to many applications, *e.g.* the diagnosis of attention deficit hyperactivity disorder (ADHD) using MRI scans. Remarkably, in this case an analytic solution to the inversion problem exists. By a star-shaped surface, we mean a parameterized surface $f \in \mathcal{F}$ that, up to a translation, can be written in the form $f(u, v) = \rho(u, v)e(u, v)$, where $\rho(u, v) \in \mathbb{R}_{\geq 0}$, and $e(u, v) \in \mathbb{S}^2$ 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 in spherical coordinates, f is given by $(r, \theta, \phi) = (\rho(u, v), u, v)$.

Note that the volume enclosed by a star-shaped surface is a star domain (*i.e.* 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). However, in addition to this purely geometric property, we also demand that the surface have a particular parameterization and that the ‘center’ be located at $r = 0$.

For this type of surface, 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

$$n^r(u, v) = \langle n(u, v), e(\theta(u, v), \phi(u, v)) \rangle, \quad (8)$$

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

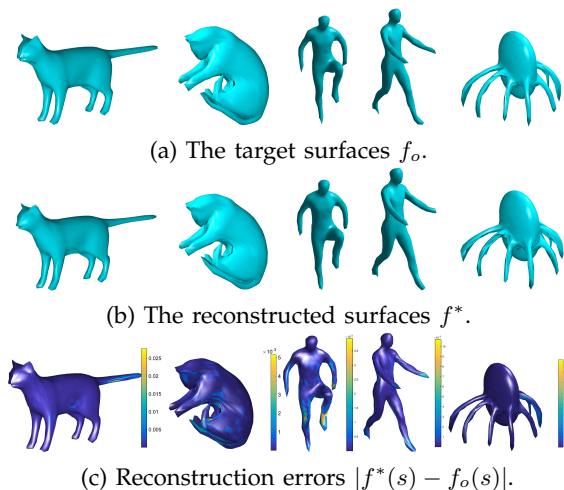


Fig. 2: Reconstruction of surfaces from their SRNF representations, using 3642 SH basis elements. Additional examples are shown in the supplementary file.

because we would not know θ and ϕ , the angular coordinates of the surface we are trying to recover. In the special parameterization, the expression becomes: $n^r(u, v) = \langle n(u, v), e(u, v) \rangle$, and this can be calculated. The result is very simple: $n^r(u, v) = \rho^2(u, v)$. 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:

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

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, 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.

2.3 Reconstruction examples

Fig. 2 shows some examples of the reconstruction of real objects. In each case, from the original surface f_o , we compute its multiresolution surface \mathbf{f}_o , compute the multiresolution SRNF $\mathbf{q}_o = Q(\mathbf{f}_o)$, and then estimate f^* by SRNF inversion of \mathbf{q}_o . We display the original surface f_o , the reconstructed surface f^* , and the reconstruction errors $|f^*(s) - f_o(s)|$. All these results were obtained using 3642 SH basis elements.

An inspection of these results shows that the multiresolution inversion procedure is able to reconstruct the original surface with a very high accuracy. For the complex articulated surfaces of Fig. 2, the iterative optimization procedure reduces the energy by **three orders of magnitude**. The plots of the reconstruction errors in Fig. 2(c) show that most of the reconstruction errors occur at high curvature regions.

TABLE 1: Computation time, in seconds, on Intel i7-3770 CPU @ 3.40Ghz with 16Gb of RAM.

	3642 Harmonic basis	100 PCA basis
Grids of 64×64	16.5 min	20.14 sec
Grids of 128×128	–	2.5 min

To quantify the impact of initialization on the convergence of the proposed algorithm, we performed an experiment in which we took 13 shape classes from the SHREC07 watertight models (see Section 4.4), with each class containing 20 models. For each model, we computed its SRNF, then re-estimated the original surface using the inversion algorithm, initialized using a unit sphere. This is a generic choice for the initialization, using no prior information about the nature of the surface, and hence is necessarily suboptimal. Finally, we measured the reconstruction error. Median reconstruction error was 0.0008 with a 25th percentile of 0.0003 and a 75th percentile of 0.0022. (All surfaces were rescaled to unit surface area.) The supplementary file includes reconstruction error analyses for each of the 13 classes of shapes.

Table 1 summarizes average computation cost for inverting one SRNF on a 3.40Ghz i7-3770 CPU with 16GB RAM. In conjunction with Table 2, it underlines the drastic improvements achieved for the statistical shape analysis tasks listed in Section 3. In conclusion, this algorithm allows us to apply statistical tools for ESA of complex surfaces for the first time.

3 DEVELOPMENT OF STATISTICAL TOOLS

The ability to invert Q enormously simplifies statistical shape modeling. In contrast with the approach in [18], where analysis is performed in \mathcal{F} under a complicated metric, the new framework performs analysis in the \mathbb{L}^2 space of SRNFs, and transforms only the end results to \mathcal{F} . 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. The basic algorithms for computing the Karcher mean shape, for parallel transport, and for transferring deformations from one shape to another, using both previous and the proposed methods, are described in Table 2. Computationally intensive steps are underlined, and the computational complexity is indicated in boxes. In the following subsections, we elaborate on the list of target analyses studied, and the mechanisms used to perform them using SRNF inversion.

In Sections 3.1, 3.2, and 3.3, we are given two surfaces f_1 and f_2 , and their SRNF representations $q_1 = Q(f_1)$ and $q_2 = Q(f_2)$. We first perform an elastic registration by solving for the optimal rotation O and optimal re-parameterization γ such that $|q_1 - O(q_2, \gamma)|$ is minimized (see [1] for details). If O^* and γ^* is

TABLE 2: Comparison of algorithms in terms of complexity and computation time.

	Previous [17], [18] (spherical grids of 32×32)	Proposed (spherical grids of 64×64)
Karcher Mean	<p><i>Algorithm 1:</i> Let μ_f^0 be an initial estimate. Set $j = 0$.</p> <ol style="list-style-type: none"> Register f_1, \dots, f_n to μ_f^j. For each $i = 1, \dots, n$, construct a geodesic to connect f_i to μ_f^j and evaluate $v_i = \exp_{\mu_f^j}^{-1}(q_i)$. Compute the average direction $\bar{v} = \frac{1}{n} \sum_{i=1}^n v_i$. If $\ \bar{v}\$ is small, stop. Else, update $\mu_f^{j+1} = \exp_{\mu_f^j}(\epsilon \bar{v})$ by shooting a geodesic, $\epsilon > 0$, small. Set $j = j + 1$ and return to Step 1. <p>n geodesics per iteration, 2 hours per geodesic.</p>	<p><i>Algorithm 2:</i> 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"> Register $q_i = Q(f_i), i = 1, \dots, n$, to \bar{q}. Update the average $\bar{q} = \frac{1}{n} \sum_{i=1}^n q_i$. 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>1 inversion, 16.5 min (with harmonic basis) or 20.14 seconds (with PCA basis).</p>
Parallel Transport	<p><i>Algorithm 3:</i> 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"> 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>1 geodesic + m parallel transports, 2 hours per geodesic.</p>	<p><i>Algorithm 4:</i> Parallel transport on \mathbb{L}^2 remains constant.</p> <ol style="list-style-type: none"> Compute $w = Q_{*,f_1}(v)$ (differential of Q). Find f by <u>inversion</u> s.t. $Q(f) = Q(f_2) + \epsilon w$, ϵ is small. Evaluate $\frac{f-f_2}{\epsilon}$ and set it to be v^{\parallel}. <p>1 inversion, 16.5 min (with harmonic basis) or 20.14 seconds (with PCA basis).</p>
Transfer Deformation	<p><i>Algorithm 5:</i></p> <ol style="list-style-type: none"> Find a geodesic $\beta(t)$ connecting f_1 to h_1 and evaluate $v = \exp_{f_1}^{-1}(h_1)$. 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"> Parallel transport $V(\frac{\tau-1}{m})$ from $\alpha(\frac{\tau-1}{m})$ to $\alpha(\frac{\tau}{m})$ and name it $V(\frac{\tau}{m})$. Shoot a geodesic $\beta'(\tau)$ from f_2 with velocity $v^{\parallel} = V(1)$ and set $h_2 = \beta'(1)$. <p>3 geodesics + m parallel transports, 2 hours per geodesic.</p>	<p><i>Algorithm 6:</i> Parallel transport on \mathbb{L}^2 remains constant.</p> <ol style="list-style-type: none"> Compute $v = Q(h_1) - Q(f_1)$. Find h_2 by <u>inversion</u> s.t. $Q(h_2) = Q(f_2) + v$. <p>1 inversion, 16.5 min (with harmonic basis) or 20.14 seconds (with PCA basis).</p>

the solution to this optimization problem, then $\bar{f}_2 = O^*(f_2 \circ \gamma^*)$ and its corresponding SRNF is given by $\bar{q}_2 = O^*(q_2, \gamma^*)$. This is equivalent to performing the analysis in the quotient space $\mathcal{S} = \mathcal{Q}/\mathcal{G}$.

3.1 Geodesic Path Reconstruction

Given two surfaces f_1 and f_2 , one wants to construct a geodesic path $\alpha(t)$, such that $\alpha(0) = f_1$ and $\alpha(1) = \bar{f}_2$. Let $\beta : [0, 1] \rightarrow \mathbb{L}^2(\mathbb{S}^2, \mathbb{R}^3)$ denote the straight line connecting q_1 and \bar{q}_2 . Then, for any arbitrary point $\beta(\tau) \in \mathbb{L}^2(\mathbb{S}^2, \mathbb{R}^3)$, we want to find a surface $\alpha(\tau) \in \mathcal{F}$ such that $\|Q(\alpha(\tau)) - \beta(\tau)\|$ is minimized. We will accomplish this using the following multiresolution analysis. Let \mathbf{f}_1 and \mathbf{f}_2 be the multiresolution representations of f_1 and \bar{f}_2 , respectively. We then compute the multiresolution SRNFs of \mathbf{f}_1 and \mathbf{f}_2 , which we denote by \mathbf{q}_1 and \mathbf{q}_2 , respectively. Recall that $\mathbf{q}_i = (q_i^1, \dots, q_i^m)$ where $q_i^j = Q(f_i^j)$. To compute the point $\alpha(\tau)$ on the geodesic path α , we compute $\alpha^j(\tau)$, for $j = 1$ to m , by SRNF inversion of $\beta^j(\tau) = (1 - \tau)q_1^j + \tau q_2^j$, using $\alpha^{j-1}(\tau)$ as an initialization for the optimization procedure. The final solution $\alpha(\tau)$ is set to be $\alpha^m(\tau)$.

For the first iteration, one can set $\alpha^0(\tau)$ to be either a unit sphere or f_1^1 . Note that Xie et al. [29] required building the geodesic path α sequentially. That is, for a small $\epsilon > 0$, Xie et al. start by solving for $\alpha(\epsilon)$, using $\alpha(0) = f_1$ as an initialization, and then use $\alpha(\epsilon)$ to initialize the next step for finding $\alpha(2\epsilon)$, and so on. This is no longer a requirement with our approach;

one can compute $\alpha(\tau)$ for any arbitrary $\tau \in [0, 1]$ without computing the entire geodesic.

Finally, if f_1 and f_2 are star-shaped surfaces, one can use the analytic solution of Section 2.2 together with numerical inversion to construct the geodesic path α as follows. First, find the geodesic in the quotient space $\mathcal{S} = \mathcal{Q}(\mathcal{F})/\mathcal{G}$ between the corresponding SRNFs, which is trivially a straight line. It is not guaranteed, however, that all the intermediate SRNFs correspond to star-shaped surfaces; thus the analytic form \tilde{f} may not be the correct inversion. One can use \tilde{f} , however, as an initial guess for the inverse, thereby better initializing the reconstruction-by-optimization problem.

3.2 Shooting Geodesics

Given a surface f and a tangent vector v_0 at f , one wants to construct a geodesic $\alpha(\tau)$ such that $\alpha(0) = f$ and $\dot{\alpha}(0) = v_0$. (Here $\dot{\alpha} = d\alpha/d\tau$.) Note that shooting a geodesic is essentially evaluating numerically the exponential map $\exp_f(\tau v_0) = \alpha(\tau), \tau \in [0, 1]$. Let β denote a straight line, i.e. $\beta(\tau) = Q(f) + \tau Q_{*,f}(v_0)$, where $Q_{*,f}$ is the differential of Q at f given by Eq. (7). Then the desired geodesic $\alpha(\tau)$ is of the form $Q(\alpha(\tau)) = \beta(\tau)$. This path $\alpha(\tau)$ is computed using the same approach as the one described in Section 3.1. Some statistical analyses computed using shooting geodesics are presented in Section 4.3.

3.3 Transferring Deformation between Shapes

Given surfaces f_1, h_1 and f_2 , we are interested in estimating the deformation from f_1 to h_1 and then




Fig. 3: Deformation transfer: 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 .

applying this deformation to f_2 . We first optimally register h_1 and f_2 onto f_1 using the approach described above. To simplify the notation, we also use f_1, h_1 and f_2 to denote the optimally aligned and reparameterized surfaces.

The deformation transfer task can be decomposed into three steps: compute the deformation v from f_1 to h_1 ; transfer v at f_1 to f_2 resulting in v^{\parallel} ; and deform f_2 into h_2 using v^{\parallel} . The first and third steps imply constructing geodesics, while the second step uses a parallel transport. With the proposed framework, h_2 can be computed by inversion of $Q(f_2) + (Q(h_1) - Q(f_1))$, which is computationally more efficient than previous work, since all the computations are done in \mathcal{Q} . The algorithm is detailed in Table 2. Fig. 3 shows an example of transferring a deformation from one surface to another using the proposed approach.

3.4 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. The mean shape μ_f is computed using Algorithm 2 of Table 2. Let $q_i, i = 1, \dots, n$ be the SRNFs of the aligned and registered surfaces in the sample; let μ_q be their average and let u_q^k be the k^{th} principal component. The k^{th} principal mode of variation for the SRNFs is given by $\mu_q \pm \lambda u_q^k, \lambda \in \mathbb{R}^+$. While these summary statistics are computed using Principal Component Analysis (PCA) in \mathcal{Q} , to visualize them, one must map them back to \mathcal{F} . Thus, the mean shape $\mu_f \in \mathcal{F}$ is computed by finding μ_f such that $Q(\mu_f) = \mu_q$. To visualize the principal directions in \mathcal{F} , we need to find f^k such that $Q(f^k) = \mu_q \pm \lambda u_q^k$. This is a geodesic shooting problem, which requires one SRNF inversion for each λ .

In terms of computational complexity, computing the mean shape only requires the inversion of one multiresolution SRNF, while previous techniques, *e.g.* [17], [18] are iterative, and required the computation of n geodesics per iteration, with n being the number of shapes to average, and each geodesic being computed using the expensive pullback metric.




Fig. 4: Linear interpolation in \mathcal{F} vs. geodesic path by SRNF inversion. The red and blue curves in (c) correspond to the energy path along the geodesic and along the linear path, respectively. The energy decreases from 10^{-2} to 10^{-4} .

3.5 Random Sampling from Shape Models

Given a sample of observed surfaces $f_1, \dots, f_n \in \mathcal{F}$, one wants to fit a probability model to the data. Estimating probability models on nonlinear spaces like \mathcal{F} is difficult: classical statistical approaches, developed for vector spaces, do not apply directly. By first computing the SRNF representations of the surfaces, one can use all the usual statistical tools in the vector space \mathcal{Q} , and then map the results back to \mathcal{F} using the proposed SRNF inversion algorithms.

Let q_1, \dots, q_n be the SRNFs of the registered surfaces f_1, \dots, f_n and $G(q)$ be the model probability distribution fitted to $\{q_1, \dots, q_n\}$. (Recall that registration and alignment is performed using the SRNF framework described above and detailed in [1].) A random sample q_s can be generated from G . We then find f_s such that $Q(f_s) = q_s$. Note that G can be an arbitrary distribution, parametric or non-parametric. In this article, we use a wrapped truncated normal distribution, which can be learned using Principal Component Analysis on \mathcal{Q} . We caution the reader that the distribution is defined and analyzed in SRNF space; a distribution on \mathcal{Q} induces a distribution on \mathcal{F} , but we have not derived it explicitly. We compute a random sample f_s for the purposes of display and validation only.

4 EXPERIMENTAL RESULTS

Sections 4.1 to 4.3 present experimental results for the shape analysis tasks described in Section 3. Section 4.4 describes two applications: the classification of generic 3D shapes, and the diagnosis of attention deficit hyperactivity disorder (ADHD) using shapes of subcortical parts in human brain.

4.1 Geodesic paths

Fig. 4 shows an example of a geodesic path between f_1 and f_2 , where f_1 is a straight cylinder and f_2 is a bent cylinder. Fig. 4(a) shows the linear path between f_1 and f_2 in \mathcal{F} (*i.e.* each pair of corresponding points is connected by a straight line). Observe how the

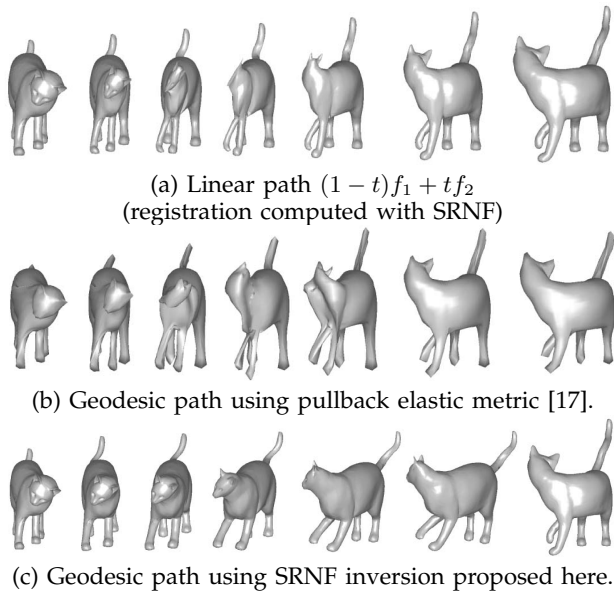


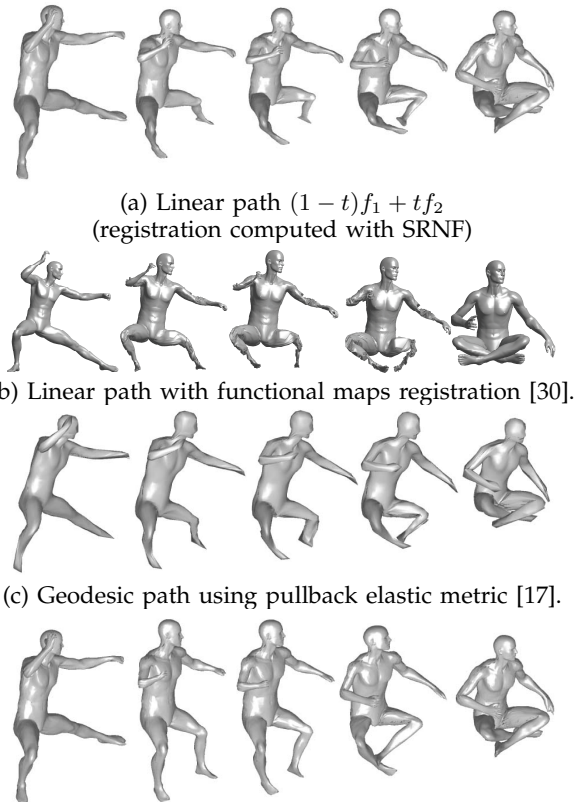
Fig. 5: Comparison between (a) linear interpolation in \mathcal{F} , (b) geodesic path estimated using the pullback metric of Kurtek et al. [17], and (c) geodesic path computed using the proposed SRNF inversion.

intermediate shapes along this path shrink unnaturally. Fig. 4(b) shows the geodesic path computed by SRNF inversion where the cylinder bends quite naturally without shrinking. Here the inversion procedure uses SH basis and is initialized with a sphere. We also plot in Fig. 4(c) the energy of Eq. (3) for each estimated surface along the geodesic path obtained by SRNF inversion (the red curve) and obtained by linear interpolation in \mathcal{F} (the blue curve). Observe that the energy of the geodesic path obtained by numerically inverting SRNFs is of order 10^{-4} , which is significantly lower than the energies of the linear path (of order 10^{-2}).

Next, we present geodesics involving highly articulated surfaces undergoing complex deformations: cat shapes undergoing isometric (*i.e.* bending) motion (Fig. 5), and human shapes in different poses undergoing both isometric and elastic deformations (Figs. 6, 7, and 8). In these examples (Figs. 5, 6, and 7), the SRNF inversion procedure uses 3642 SH basis elements, while for Fig. 8, inversion uses only 100 PCA basis elements computed from a collection of 398 human shapes. We also compare our results with:

1. **Linear interpolation in \mathcal{F} with SRNF Registration:** Despite good quality registration between shapes, this approach leads to unnatural deformations, as shown in Figs. 5(a), 6(a), 7(a), and 8(a).

2. **Linear interpolation in \mathcal{F} with Functional Map Registration [30]:** As shown in Fig. 6(b), functional maps do not produce correct one-to-one correspondences. As a result, the deformation between surfaces contains many artifacts and degeneracies. In addition, functional maps are: (1) limited to isometric surfaces,



(d) Geodesic path using SRNF inversion proposed here.

Fig. 6: Comparison of our approach with different frameworks. Observe that the in-between shapes in (b) contain artifacts due to the inaccurate correspondences computed using functional maps.

i.e. surfaces that only bend; (2) require pre-segmented shapes; and (3) require part-wise correspondences as input [30]. Our approach is fully automatic, and finds both registrations and deformations between surfaces that undergo large bending and stretching.

3. **Pullback Metric of Kurtek et al. [17]:** This approach is computationally very expensive, in particular when computing summary statistics, see the discussion of Table 2. The underlying metric is not translation invariant and can be trapped in local minima when dealing with complex surfaces: see Figs. 5(b) and 6(c).

In all examples, our approach produces better results. In particular, when the deformations between the source and target surfaces are significant, the intermediate shapes produced by linear interpolation in \mathcal{F} contain self-intersections and unnatural shrinkages.

4.2 Geodesic shooting and deformation transfer

Fig. 9 shows several examples of deformation transfer across human body shapes using geodesic shooting and parallel transport (Sections 3.2 and 3.3). The source deformations, *i.e.* the surfaces f_1 and h_1 , are shown in Fig. 9(a) while Fig. 9(c) shows these deformations transferred onto another subject using

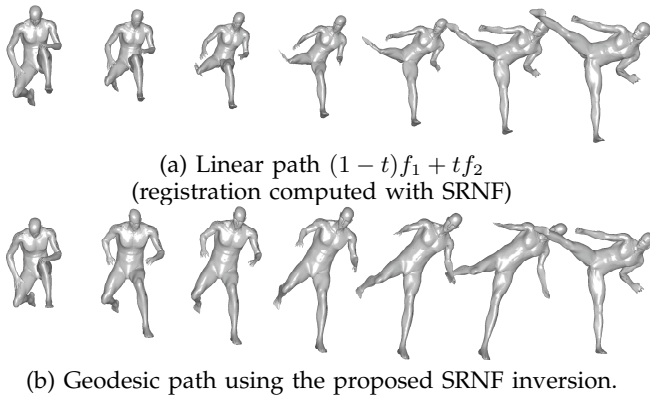


Fig. 7: Comparison between linear interpolations in \mathcal{F} and geodesic paths by SRNF inversion.

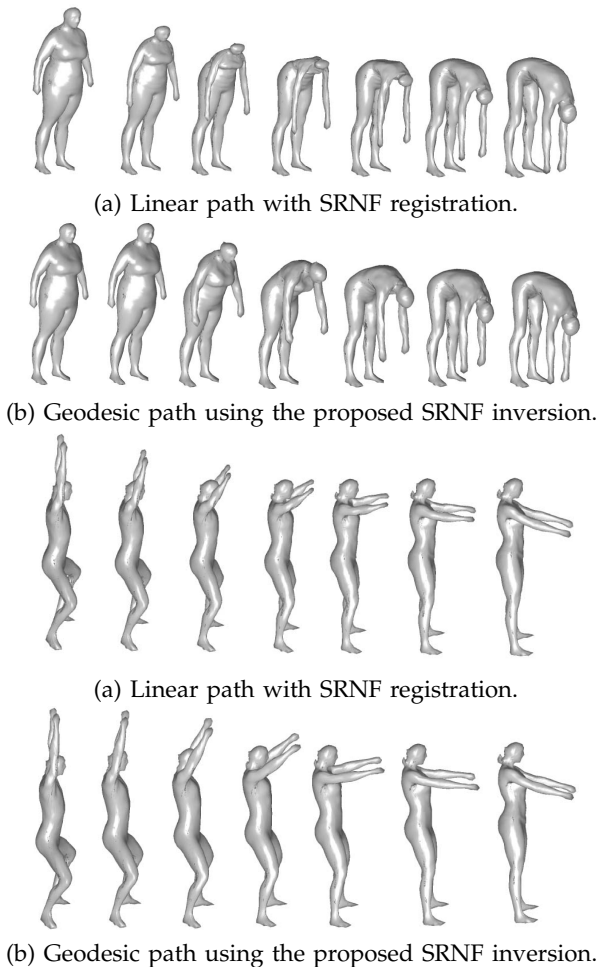


Fig. 8: Comparison between linear interpolations in \mathcal{F} and geodesic paths by SRNF inversion (obtained using 100 PCA basis elements).

geodesics obtained by SRNF inversion. For comparison, we also show in Fig. 9(b) the deformations transferred by linear extrapolation in \mathcal{F} . Observe that the deformations transferred by SRNF inversion look very natural and are free from artifacts, self-intersections, and unnatural elongations, contrary to

the deformations obtained by linear extrapolation.

4.3 Statistical analysis of surfaces

We consider a collection of 398 human shapes [20] composed of multiple subjects in different poses. The first subject has 35 different poses including a neutral one. All other subjects have a neutral pose and a few other poses. We first compute the SRNF representations of these surfaces, perform statistical analysis in the quotient space, $\mathcal{S} = \mathcal{Q}/\mathcal{G}$, using standard linear statistics such as Principal Component Analysis, and finally map the results to the surface space, \mathcal{F} , using the proposed SRNF inversion algorithm.

Fig. 10(a) shows the mean shape and the first three modes of variation computed using all the subjects in a neutral position. We refer to this statistical model as *pose-independent*. Similarly, we consider the 35 different poses of the first subject and compute their mean shape and modes of variation (Fig. 10(b)). We call the resulting model the *pose shape model*.

We use these two statistical models to synthesize random human shapes in arbitrary poses. Let f_1 be the neutral pose in the 35 models we used to learn the pose shape model. First, we randomly generate a surface f_2 that is in a neutral pose using the pose-independent statistical model. We then generate an arbitrary pose h_1 using the pose shape model. Finally, we use the deformation transfer procedure (Section 3.3) to deform f_2 in the same way that f_1 deforms into h_1 . This results in a new random human shape h_2 in an arbitrary pose. Fig. 10(c) shows 20 arbitrary human shapes generated with this method. This approach uses two separate models to capture inter- and within-subject variabilities; used with the deformation transfer mechanism, this provides greater control over the generation of different subjects and poses. One could also consider all the human shapes together irrespective of their pose and learn a single statistical model from which random instances can be sampled.

Table 2 summarizes the computational benefits of our framework compared to previous techniques. In [17], computing the mean shape of N objects would require $N \times m$ geodesics, where m is the number of iterations in their iterative algorithm. This is very expensive. This is also true for methods such as [10], [31], which have a similar approach for geodesic construction. In addition, [10], [31] assume that the input surfaces are fully registered, using another, unrelated, method. Thus, these prior approaches, based on finding individual geodesics using the pullback metric in \mathcal{F} , do not scale to large datasets. Our approach simultaneously finds correspondences and geodesics. Its computational cost is limited to the cost of finding a Euclidean mean and inverting one SRNF. Thus, it is computationally very efficient and also scalable.

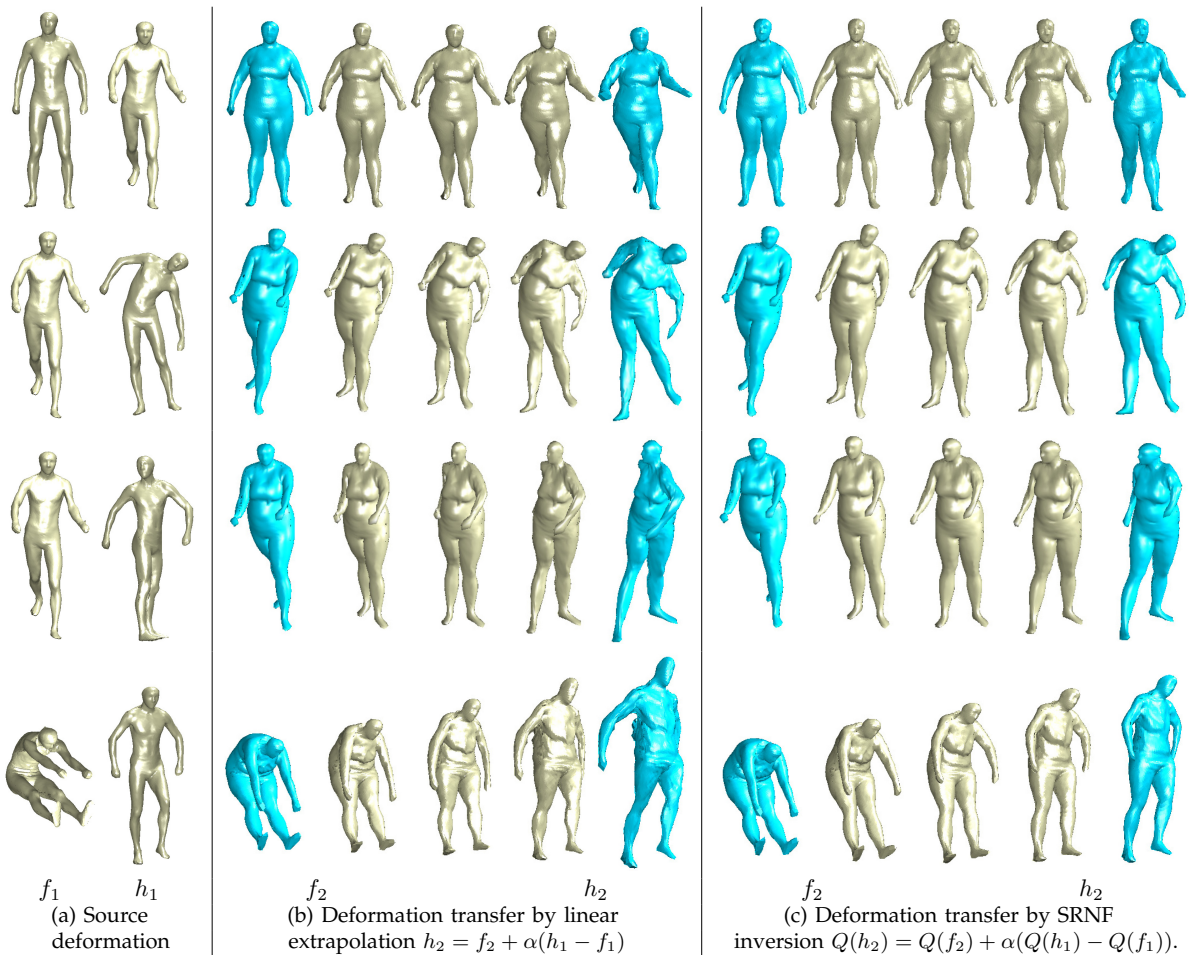


Fig. 9: Examples of deformation transfer using geodesic shooting in the space of SRNFs. The middle shapes, in gray, are intermediate shapes along the deformation paths.

TABLE 3: Classification performance, in %, on the genus-0 watertight 3D models of SHREC07 [19].

	Gauss SRNF	Gauss \mathcal{F}	1NN SRNF	3NN SRNF	1NN \mathcal{F}	3NN \mathcal{F}
w/o regist.	52.42	44.49	52.86	50.66	56.38	52.42
ICP [32]	–	–	–	–	82.81	82.37
Elastic regist.	96.47	91.18	100	99.56	96.91	96.91
Other methods	BoC [33]	Spatial BoC [33]	Hybrid BoW [34]	Tabia et al. [35]		
1NN	93.25	92.5	91.8	85.3		

4.4 Classification Applications

Generic 3D shape classification. We apply the proposed approach to the important problem of classifying generic 3D shapes. We use the SHREC07 watertight 3D model benchmark [19], which is composed of 400 watertight 3D models evenly divided into 20 shape classes. We only consider the 13 classes that are composed of genus-0 surfaces. First, we compute spherical parameterizations [36]; then we map them to \mathcal{Q} . To classify a shape, we use:

- Leave-One-Out k -nearest neighbor classifier (kNN-SRNF) using the \mathbb{L}^2 metric defined on the shape space $\mathcal{S} = Q(\mathcal{F})/\mathcal{G}$, i.e. the \mathbb{L}^2 metric between the aligned and registered SRNFs.
- The Gaussian classifier (Gauss-SRNF) on principal components, in the space of SRNFs as described in Section 3. We use the Leave-One-Out (LOO) strategy. We classify each query to the shape category whose mean is the closest to the query in terms of Mahalanobis distance.

We also compare our results to those obtained using:

- The Gaussian classifier and the kNN classifier, but in the surface space \mathcal{F} , with and without elastic registration (using the proposed metric);
- The iterative closest point (ICP) algorithm [32];
- The covariance descriptors [33], the hybrid Bag of Words [34] and the approach of [35].

Table 3 reports the classification rates, and shows that the best performance is attained by the proposed kNN-SRNF and Gauss-SRNF classifiers. The results suggest that the parametrization-invariant metric and probability models in \mathcal{Q} result in better 3D shape classification.




Fig. 10: Statistical analysis of 3D human body shapes using the SRNF inversion proposed in this article.

ADHD classification. Finally, we apply our methods to shape-based diagnosis of attention deficit hyperactivity disorder (ADHD) using MRI scans. The surfaces of certain subcortical brain structures used here were extracted from T1 weighted brain MR images of young adults aged between 18 and 21. These subjects were recruited from the Detroit Fetal Alcohol and Drug Exposure Cohort [37], [38]. Among the 34 subjects, 19 were diagnosed with ADHD and the remaining 15 were controls (non-ADHD). (Some examples of left structures are shown in Fig. 5 of the supplementary material.) First we register the extracted surfaces as described in [1], then map them into \mathcal{Q} using Q . In order to distinguish ADHD and control samples, we use the Gaussian classifier on principal components \mathcal{Q} as described in Section 3.

Table 4 shows the LOO nearest neighbor classification rate of our approach compared to five other state-of-the-art techniques [15], [32], [39]. The best performance is attained using the proposed SRNF Gaussian classifier between left putamen surfaces. These results suggest that the parametrization-invariant metric and the probability models in our approach provide 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

TABLE 4: Classification performance, in %, for six different techniques.

	SRNF Gauss Proposed	SRM Gauss [39]	SRM NN [40]	Har- monic	ICP [32]	SPHARM PDF [15]
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

related ideas—it also provides significantly-improved classification performance.

5 CONCLUSIONS

The SRNF representation is potentially an important tool in the statistical shape analysis of surfaces. Since SRNF space is a vector space with the \mathbb{L}^2 -metric, statistical analysis in this space is simple and low-cost compared to previous methods that perform analysis directly in surface space. Use of the SRNF has been limited, however, due to lack of a method for reconstructing the corresponding surface. We have introduced methods for approximating this inverse mapping, thereby removing the obstacle. As a result, by adopting the proposed framework, the computational cost of algorithms for the statistical analysis of surface shape can be reduced by an order of magnitude. Experimental results show the use of the proposed framework in various 3D shape analysis tasks, such as computing geodesics, transferring deformations, computing statistical shape summaries, and fitting and sampling from probabilistic models. We have also demonstrated that the method achieves state-of-the-art performance in the classification of generic 3D shapes, and in the analysis of brain structures arising from ADHD data.

Acknowledgment. The authors would like to thank: Nils Hasler for providing us with the 3D human shape models; Sebastian Kurtek for sharing programs for registration of SRNFs under the \mathbb{L}^2 metric; Raif Rustamov and Maks Ovsjkanikov for the discussion about functional maps and for sharing their data and results. AS gratefully acknowledges support from a Durham University IAS Senior Research Fellowship (EU grant FP7-609412) and the DU Department of Mathematical Sciences.

REFERENCES

- [1] I. Jermyn *et al.*, “Elastic shape matching of parameterized surfaces using square root normal fields,” *Computer Vision, IEEE European Conference on*, vol. 5, no. 14, pp. 805–817, 2012.
- [2] I. Dryden and K. Mardia, *Statistical Shape Analysis*. John Wiley & Son, 1998.
- [3] B. Allen *et al.*, “The space of human body shapes: Reconstruction and parameterization from range scans,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 587–594, Jul. 2003.
- [4] S. Bouix *et al.*, “Hippocampal shape analysis using medial surfaces,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2001, vol. 2208, pp. 33–40.

- [5] K. Gorczowski *et al.*, "Multi-object analysis of volume, pose, and shape using statistical discrimination," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 32, no. 4, pp. 652–61, 2010.
- [6] S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces (Applied Mathematical Sciences)*, 2003rd ed., 2002.
- [7] R. Davies *et al.*, "Building 3-d statistical shape models by direct optimization," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 4, pp. 961–981, April 2010.
- [8] M. Kilian *et al.*, "Geometric modeling in shape space," in *ACM SIGGRAPH*, 2007.
- [9] B. Heeren *et al.*, "Time-discrete geodesics in the space of shells," *CGF*, vol. 31, no. 5, pp. 1755–1764, 2012.
- [10] C. Zhang *et al.*, "Shell pca: statistical shape modelling in shell space," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1671–1679.
- [11] O. Van Kaick *et al.*, "A Survey on Shape Correspondence," *CGF*, vol. 30, no. 6, pp. 1681–1707, 2011.
- [12] H. Zhang *et al.*, "Deformation-driven shape correspondence," *CGF*, vol. 27, no. 5, pp. 1431–1439, 2008.
- [13] T. Windheuser *et al.*, "Geometrically consistent elastic matching of 3d shapes: A linear programming solution," in *Computer Vision, IEEE International Conference on*, 2011, pp. 2134–2141.
- [14] C. Brechbühler *et al.*, "Parameterization of closed surfaces for 3D shape description," *Computer Vision and Image Understanding*, vol. 61, no. 2, pp. 154–170, 1995.
- [15] M. Styner *et al.*, "Framework for the statistical shape analysis of brain structures using spharm-pdm," *Insight Journal*, 2006.
- [16] A. Srivastava *et al.*, "Shape analysis of elastic curves in euclidean spaces," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 7, pp. 1415–1428, July 2011.
- [17] S. Kurtek *et al.*, "Elastic geodesic paths in shape space of parameterized surfaces," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 9, pp. 1717–1730, 2012.
- [18] Q. Xie *et al.*, "Parallel transport of deformations in shape space of elastic surfaces," in *Computer Vision, IEEE International Conference on*, December 2013.
- [19] D. Giorgi *et al.*, "Shape retrieval contest 2007: Watertight models track," *SHREC competition*, vol. 8, 2007.
- [20] N. Hasler *et al.*, "A statistical model of human pose and body shape," in *CGF*, vol. 28, no. 2, 2009, pp. 337–346.
- [21] P. Michor and R. Bryant. Representing immersions from a surface into 3-space. [Online]. Available: <https://mathoverflow.net/questions/128441/representing-immersions-from-a-surface-into-3-space>
- [22] K. Abe and J. Erbacher, "Isometric immersions with the same Gauss map," *Mathematische Annalen*, vol. 215, no. 3, pp. 197–201, 1975.
- [23] J.-H. Eschenburg *et al.*, "Compatibility of Gauss maps with metrics," *Differential Geometry and its Applications*, vol. 28, no. 2, pp. 228–235, 2010.
- [24] D. Hoffman and R. Osserman, "The Gauss map of surfaces in \mathbb{R}^3 and \mathbb{R}^4 ," *Proceedings of the London Mathematical Society*, vol. 3, no. 1, pp. 27–56, 1985.
- [25] V. I. Arnold, "Ten problems," *Adv. Soviet Math*, vol. 1, pp. 1–8, 1990.
- [26] —, *Arnold's Problems*. Springer Berlin Heidelberg, 2004.
- [27] E. Klassen, private communication.
- [28] B. T. T. Yeo *et al.*, "On the construction of invertible filter banks on the 2-sphere," *Image Processing, IEEE Transactions on*, vol. 17, no. 3, pp. 283–300, 2008.
- [29] Q. Xie *et al.*, "Numerical inversion of srnfs for efficient elastic shape analysis of star-shaped objects," in *Computer Vision—ECCV 2014*, 2014, pp. 485–499.
- [30] M. Ovsjanikov *et al.*, "Functional maps: a flexible representation of maps between shapes," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, p. 30, 2012.
- [31] B. Berkels *et al.*, "Discrete geodesic regression in shape space," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*. Springer, 2013, pp. 108–122.
- [32] P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 14, no. 2, pp. 239–256, Feb 1992.
- [33] H. Tabia *et al.*, "Covariance descriptors for 3d shape matching and retrieval," in *Computer Vision and Pattern Recognition, IEEE Conference on*. IEEE, 2014, pp. 4185–4192.
- [34] G. Lavoué, "Combination of bag-of-words descriptors for robust partial shape retrieval," *The Visual Computer*, vol. 28, no. 9, pp. 931–942, 2012.
- [35] H. Tabia *et al.*, "A new 3d-matching method of nonrigid and partially similar models using curve analysis," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 4, pp. 852–858, 2011.
- [36] S. Kurtek *et al.*, "Landmark-guided elastic shape analysis of spherically-parameterized surfaces," *CGF*, vol. 32, no. 2pt4, pp. 429–438, 2013.
- [37] S. W. Jacobson *et al.*, "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*, vol. 28, no. 11, pp. 1732–1745, 2004.
- [38] M. J. Burden *et al.*, "An event-related potential study of response inhibition in adhd with and without prenatal alcohol exposure," *Alcoholism: Clinical and Experimental Research*, vol. 34, no. 4, pp. 617–627, 2010.
- [39] S. Kurtek *et al.*, "Parameterization-invariant shape statistics and probabilistic classification of anatomical surfaces," in *IPMI*, 2011, vol. 6801, pp. 147–158.
- [40] —, "Parameterization-invariant shape comparisons of anatomical surfaces," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 3, pp. 849–858, 2011.



Hamid said received the PhD degrees in Computer Science from Tokyo Institute of Technology in 2006. He is currently Associate Professor at Murdoch University (Australia) and Adjunct Associate Professor with the Phenomics and Bioinformatics Research Centre (PBRC) of the University of South Australia (UniSA). His research interests span various fields of computer vision, computer graphics, and pattern recognition, with a special focus on the 3D acquisition, modeling and analysis of static and deformable 3D objects.



Qian Xie received her bachelor degree from Renmin University, China in 2007, and master degree in Statistics from American University in 2009. After that, she obtained the PhD degree in Statistics from the Department of Statistics at Florida State University in 2015. Her main research interests include statistical shape analysis, image analysis and their applications to biological and medical problems.



Ian H. Jermyn is Reader in Statistics in the Department of Mathematical Sciences at Durham University. His research concerns 'statistical geometry': the statistical modelling of shape and geometric structure, particularly using random fields with long-range and higher-order interactions, and Riemannian geometry. This work is motivated by, and has been extensively applied to, problems in image processing and computer vision, and in particular remote sensing imagery.



Anuj Srivastava is a Professor of Statistics at the Florida State University in Tallahassee, FL. He obtained his MS and PhD degrees in Electrical Engineering from the Washington University in St. Louis in 1993 and 1996, respectively. After spending the year 1996-97 at the Brown University as a visiting researcher, he joined FSU as an Assistant Professor in 1997. His research is focused on pattern theoretic approaches to problems in image analysis, computer vision, and signal processing. Specifically, he has developed computational tools for performing statistical inferences on certain nonlinear manifolds. He is a Fellow of IEEE and IAPR.