

Elastic 3D Shape Analysis Using Square-Root Normal Field Representation

Hamid Laga¹, Ian H. Jermyn², Sebastian Kurtek³ and Anuj Srivastava⁴

Abstract—Shape is an important physical property of natural and man-made 3D objects that characterizes their external appearances. Understanding differences between shapes, and modeling the variability within and across shape classes, hereinafter referred to as *shape analysis*, are problems fundamental to many applications, ranging from computer vision and computer graphics to biology and medicine. This paper provides an overview of some of the recent techniques for studying the shape of 3D objects that undergo non-rigid deformations including bending and stretching. We will mainly focus on a new representation called the square-root normal field (SRNF), discuss its properties, and show its application in the analysis of the shape of various types of objects, including human body shapes, anatomical organs such as carpal bones, and hand-drawn 2D sketches. We will show how the representation is used for (1) jointly computing correspondences and geodesics; (2) computing summary statistics such as means and modes of variations; and (3) exploring shape variability in a collection of 3D objects.

I. INTRODUCTION

Shape is an important physical property of natural and man-made 3D objects that characterizes their external appearances. Comparing the shape of objects and modelling shape variability within and across object categories are fundamental problems and building blocks in many applications. In medicine, for example, many diseases can be linked to alterations in the shape of anatomical organs [1], [2]. 3D modelling tools in computer graphics and computer-aided design are increasingly relying on the analysis of existing large collections of 3D models for creating new variations with rich geometric, structural, and aesthetic features. In computer vision, shape priors, learned from collections of 3D models, are often used to constrain ill-posed problems such as the 3D reconstruction of deformable objects from a few images or noisy range scans.

Given its importance, shape has been studied for decades by researchers from disciplines as various as mathematics, statistics, biology, computer vision, and computer graphics. As such, the terms *shape* and *shape analysis* may have different meanings and refer to different tasks. In this paper, we use shape to indicate the property of an object that characterizes its external geometry after variations due to translations, scaling, and rotations have been factored out.

This definition, introduced by D. Kendall [3], is the most commonly adopted by various communities. Similarly, by shape analysis, we refer to mathematical and algorithmic tools for:

- **Comparing the shapes of 3D objects**, *i.e.* saying whether two objects have similar shapes or not, and more importantly, quantifying and localizing the similarities and differences;
- **Computing summary statistics and prototypes**, such as the mean shape, modes of variation, and higher order statistics, of a collection of 3D models;
- **Mathematical modeling of shape variations**, using, for example, probability distributions as generative models; these models form priors for random sampling and for statistical inferences; and
- **Exploring shape variations** for synthesizing valid shapes and performing interpolations, extrapolations, statistical inferences, and regressions.

Implementing these ideas, for 3D models, requires solving several challenges, each one of which has been the subject of important research and contributions. The first concerns the representation of shape. When selecting any representation, it is important to ensure that two different shapes cannot have the same representation and that a given representation can always be associated with a valid shape. Second, almost any shape analysis task requires some measure of dissimilarity, hereinafter referred to as a *metric*. It is used to quantify deformations that affect shape (e.g. bending and stretching) while being invariant to transformations such as translation, scaling, rotation, and re-parameterization, which preserve shape. The third difficulty in shape analysis is registration, *i.e.* matching points across objects. Any shape metric requires a registration component to help decide which point on one object is compared to which point on the other. This is often a difficult problem to solve, especially between objects differing by large elastic deformations and pose variability. Thus, it is not surprising that a large body of literature is specifically dedicated to this problem; see for example [4] for a detailed survey of this topic.

Several approaches and frameworks have been proposed in the literature for solving these fundamental problems [5], [6], [7]. A prominent statistical shape analysis framework, pioneered by Kendall's school [8], [9], works with point sets that are already registered, and focuses only on deformations. Approaches such as medial surfaces [10], [11] and level sets [12] either presume registration, or solve for it using some independent pre-processing criterion such as

¹Hamid Laga is with the School of Engineering and IT, Murdoch University (Australia). H.Laga@murdoch.edu.au

²Ian H. Jermyn is with the Department of Mathematical Sciences, Durham University (UK). i.h.jermyn@durham.ac.uk

³Sebastian Kurtek is with the Department of Statistics, The Ohio State University (USA). kurtek.1@stat.osu.edu

⁴Anuj Srivastava is with the Department of Statistics, Florida State University (USA). anuj@stat.fsu.edu

MDL [13]. Kilian *et al.* [14] represent surfaces by discrete triangulated meshes and compute geodesic paths (deformations) between them, while assuming that the meshes are registered. Heeren *et al.* [15] proposed a method for computing geodesic-based deformations of thin shell shapes, with extensions for computing summary statistics in the shell space [16], but with known registration. Some other papers solve for registration [4], [17] while ignoring deformation. Examples include Windheuser *et al.* [18], 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 [19], [20] seek uniform sampling on the domain to address the registration issue. This is a major restriction since it limits registration of corresponding features across surfaces.

A majority of these papers 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.* [21] analyzed 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.

In this paper, we focus on a recent mathematical framework for elastic shape analysis, termed square-root normal fields (SRNFs). The representation was initially developed in [22], and then advanced in many other papers [23], [24], [25]. We will review the properties of this representation and demonstrate its application in many elastic shape analysis tasks, including correspondence and registration, classification and clustering, computation of statistical summaries, and shape variability exploration. While the representation has been introduced for the analysis of genus-0 surfaces, we also show that it can be used for the analysis of the shape of planar objects that undergo topological deformations.

II. MATHEMATICAL FRAMEWORK

We represent the boundary of a 3D object as a function $f : D \rightarrow \mathbb{R}^3$, which assigns to every point $s \in D$ a three-dimensional point $f(s) = (x(s), y(s), z(s))$. The choice of the domain D depends on the application. For instance, open surfaces such as 3D human faces can be embedded on a planar domain $D = [0, 1]^2$. For closed genus-0 surfaces, the unit sphere $D = \mathbb{S}^2$ is the most natural choice. Let \mathcal{F} denote the space of all such surfaces.

A. Pre-shape and shape spaces

An important challenge, when designing a framework for shape analysis, is to discount effects or variables that do not affect the shape of an object and to account only for those that do affect shape. Kendall [3] defined shape as the property of an object that remains once rotation, translation, and scale are removed from its representation. As a result, in the field of shape analysis these transformations are termed

shape-preserving; they are nuisance variables that should be discarded.

Translation and scaling are probably the easiest to deal with. One can discard translation by first translating f so that its center of mass is located at the origin: $f(s) \rightarrow f(s) - \frac{\int_D a(s)f(s)}{\int_D a(s)}$. Here $a(s)$ is the local surface area at s . It is defined as the norm of the normal vector to the surface at s . In terms of a surface parameterization, $s = (u, v)$, the normal vector to the surface at s is $n(s) = \frac{\partial f}{\partial u} \times \frac{\partial f}{\partial v}$. The scale component can be also discarded by normalizing the surface f in such a way that it has unit surface area: $f(s) \rightarrow \frac{f(s)}{\sqrt{\int_D a(s)}}$. The space \mathcal{C}_f of all origin-centered and scale-normalized surfaces is called *pre-shape space*.

In addition to scale and translation, 3D objects undergo two other shape-preserving transformations: rotations and re-parameterizations. Rotations, denoted by O , are elements of $SO(3)$, the 3×3 orthogonal matrices. They transform each point $f(s) \in \mathbb{R}^3$ into $Of(s)$. Here, we denote by Of the rotated version of the surface f . Re-parameterization is a diffeomorphism $\gamma : D \rightarrow D$, which transforms a surface f into $f \circ \gamma$. Let Γ denote the space of all such diffeomorphisms. Diffeomorphisms are shape preserving transformations; that is, the surfaces f and $f \circ \gamma$ have the same shape, albeit with different parameterizations. Note that re-parameterization is important in 3D shape analysis because it provides registration.

Invariance to rotations and re-parameterizations can be dealt with algebraically. The idea, in the case of rotations for example, is that a surface f and any other surface obtained by rotating f have the same shape. That is, $\forall O \in SO(3)$, f and Of are equivalent. Similarly, a surface f and any of its possible re-parameterizations $f \circ \gamma$ are equivalent. Thus the set $[f] = \{Of \circ \gamma, O \in SO(3), \gamma \in \Gamma\}$ forms an equivalence class under the action of the rotation and re-parameterization groups. The set $\mathcal{S}_f = \mathcal{C}_f / SO(3) \times \Gamma$ of all equivalence classes is called *shape space*.

B. Deformation-driven shape analysis

We are given two origin-centered and scale-normalized surfaces $f_1, f_2 \in \mathcal{C}_f$ and we seek to quantify the difference between their shapes. This can be done by defining a measure of dissimilarity $d(\cdot, \cdot)$ which quantifies deformations that change shape while remaining invariant to those that do not affect shape. That is:

$$d(f_1, f_2) = d_S([f_1], [f_2]) = \min_{O, \gamma} d_C(f_1, Of_2 \circ \gamma), \quad (1)$$

where d_C is a measure of dissimilarity on \mathcal{C} . The minimization in equation (1) is over all possible rotations $O \in SO(3)$ and re-parameterizations $\gamma \in \Gamma$ of f_2 . Thus, registration is embedded in the process of comparing the shape of the two surfaces f_1 and f_2 .

In shape analysis, a dissimilarity measure quantifies the amount of deformation, or energy, that one needs to apply to one shape in order to align it to the other. Consider the example of Fig. 1, where f_1 is a straight cylinder and f_2 is a bended one. The deformation of f_1 onto f_2 results in

a sequence of m intermediate shapes $F_i : \mathbb{S}^2 \rightarrow \mathbb{R}^3, i = 1, \dots, m$ such that $F_0 = f_1, F_m = f_2$, and $F_i = F_{i-1} + V_{i-1}$. Here, $V_{i-1} : D \rightarrow \mathbb{R}^3$ is a (deformation) vector field such that $F_i(s) = F_{i-1}(s) + V_{i-1}(s)$. The sequence $F = \{F_i, i = 1, \dots, m\}$ can be seen as a path or a curve in the preshape space \mathcal{C}_f . When m is sufficiently large, F can be seen as a parameterized path $F : [0, 1] \rightarrow \mathcal{C}_f$ such that $F(0) = f_1, F(1) = f_2$. Its length is given by:

$$d_{\mathcal{C}}(f_1, f_2) = L(F) = \int_0^1 \left\langle \left\langle \frac{dF}{d\tau}(\tau), \frac{dF}{d\tau}(\tau) \right\rangle \right\rangle^{\frac{1}{2}} d\tau. \quad (2)$$

Here, $\frac{dF}{d\tau}(\tau)$ is an infinitesimal vector field that deforms the surface $F(\tau)$. It can be interpreted as a small perturbation of the surface. The inner product, $\langle \langle \cdot, \cdot \rangle \rangle$, also known as *the metric*, measures the strength of this vector field. Thus, integrating over τ provides the length of the path.

In general, a surface f is treated as an element, or a point, in a pre-shape space \mathcal{C}_f . A vector field that deforms f is then a vector δf that is tangent to \mathcal{C}_f at f . Let $T_f(\mathcal{C}_f)$ denote the tangent space to \mathcal{C}_f at f . A metric is an inner product on $T_f(\mathcal{C}_f)$. It takes two tangent vectors v_1 and $v_2 \in T_f(\mathcal{C}_f)$ and returns their inner product $\langle \langle v_1, v_2 \rangle \rangle$. The norm of a vector $v \in T_f(\mathcal{C}_f)$, according to the metric $\langle \langle \cdot, \cdot \rangle \rangle$, is given by $\langle \langle v, v \rangle \rangle^{\frac{1}{2}}$. Since there are many paths that deform f_1 onto f_2 , we are interested in the shortest one, under the metric, which is called *the geodesic* in the pre-shape space \mathcal{C}_f :

$$F^* = \arg \min_F L(F). \quad (3)$$

The length of the geodesic path is the dissimilarity, in the preshape space, between the two surfaces:

$$d_{\mathcal{C}}(f_1, f_2) = \inf L(F) = L(F^*). \quad (4)$$

Combining Equations (4), (3), and (1), we obtain the dissimilarity in shape space between two surfaces f_1 and f_2 .

Under this formulation, the registration problem becomes a problem of finding the optimal rotation O^* and diffeomorphism γ^* such that:

$$(O^*, \gamma^*) = \operatorname{argmin}_{O \in SO(3), \gamma \in \Gamma} d_{\mathcal{C}}(f_1, O f_2 \circ \gamma). \quad (5)$$

The geodesic path is then the shortest path, with respect to the metric, between f_1 and $O^* f_2 \circ \gamma^*$.

C. Choice of the metric

Since almost every 3D shape analysis task includes a step in which shapes are compared based on some measure of dissimilarity, the choice of the metric is very critical for the subsequent analysis tasks. A simple and popular choice is the Euclidean distance, or \mathbb{L}^2 metric, in \mathcal{C}_f , which has been extensively used in a wide range of applications, e.g. archaeology, astronomy, morphometrics, medical diagnosis [26], [27], [28], [29]. The Euclidean distance in \mathcal{C}_f , however, has several limitations. First, it does not capture natural deformations. This is best illustrated with the example of Figs. 1-(a) and (b), which show optimal deformation paths between two surfaces f_1 and f_2 , where f_1 is a straight cylinder and f_2 is a bent one. Fig. 1-(a) shows the linear path

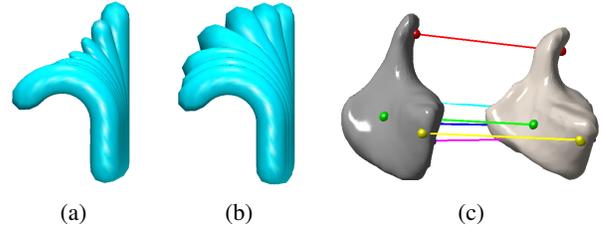


Fig. 1. Linear interpolation in \mathcal{F} vs. geodesic path by SRNF inversion. (a) Linear path $((1 - \tau)f_1 + \tau f_2)$, (b) Geodesic path by SRNF inversion, and (c) example of landmark correspondences computed with the elastic shape analysis framework.

between f_1 and f_2 (after full registration) obtained using the \mathbb{L}^2 metric by connecting each pair of corresponding points with a straight line. The intermediate shapes along this path shrink unnaturally. Fig. 1-(b) on the other hand shows a geodesic computed using a proper metric.

The second limitation of the \mathbb{L}^2 metric on \mathcal{C}_f is that the action of Γ on \mathcal{C}_f is not by isometry, *i.e.* in general, $\|f_1 - f_2\| \neq \|f_1 \circ \gamma - f_2 \circ \gamma\|$ for $\gamma \in \Gamma$. Thus, the \mathbb{L}^2 metric is not suitable for re-parameterization invariant shape analysis.

Instead of using the \mathbb{L}^2 metric on \mathcal{C}_f , one would like to explicitly capture and quantify the deformations that one needs to apply to f_1 in order to align it onto f_2 . Such deformations can be of two types: bending and stretching. This requires redefining $d_{\mathcal{C}}$ of Eqn. (1) in terms of an energy function that penalizes these deformations. There are two important pieces of work that implemented this model. The first one is inspired by the elasticity theory in physics [30] where surfaces are treated as *thin shells*, *i.e.* a thin three-dimensional material of thickness δ . Stretching in this case is caused by in-layer (tangential) shear or compression while bending is caused by friction due to transversal shear [15].

The second one showed that with specific choices of the bending and stretching terms as well as their weights, the deformation model reduces to an \mathbb{L}^2 metric in the space of square-root normal fields (SRNF), a special representation of surfaces [22]. Here, we will focus on this representation since it offers several computational benefits as discussed in the following sections.

D. SRNF representation

Jermyn et al. [22] introduced a new representation of surfaces called the *square-root normal field* (SRNF). This is essentially the field of surface normals scaled by the square root of the local area:

$$Q(f)(s) = q(s) = \frac{\mathbf{n}(s)}{\sqrt{|\mathbf{n}(s)|}}, \quad (6)$$

where $\mathbf{n}(s) = \frac{\partial f}{\partial u} \times \frac{\partial f}{\partial v}$ and $s = (u, v) \in \mathbb{S}^2$. The space of SRNFs, hereinafter denoted by \mathcal{Q} , has very nice properties that are relevant to shape analysis. In particular, Jermyn et al. [22] showed that the \mathbb{L}^2 metric on \mathcal{Q} is a special case of a full elastic metric that quantifies shape differences as a weighted sum of bending and a stretching components. Let

δq be a tangent vector to \mathcal{Q} at $q = Q(f)$. Then:

$$\langle\langle \delta q, \delta q \rangle\rangle = \int_{\mathbb{S}^2} ds \sqrt{|g(s)|} \left\{ \frac{1}{16} (g(s)^{-1} \delta g(s))^2 + \langle \delta \tilde{n}(s), \delta \tilde{n}(s) \rangle \right\}. \quad (7)$$

Here $\langle \cdot, \cdot \rangle$ denotes the standard inner product in \mathbb{R}^3 ; $g(s)$ is the determinant at s of the metric induced on the surface (also known as the first fundamental form); and $\tilde{n}(s)$ is the unit normal vector of the surface at s . Since $\sqrt{|g|} \equiv |n|$, the Eqn. (7) is equivalent to

$$\langle\langle \delta q, \delta q \rangle\rangle = \int_{\mathbb{S}^2} ds \left\{ \frac{1}{4} \frac{\delta a(s) \delta a(s)}{a(s)} + a(s) \langle \delta \tilde{n}(s), \delta \tilde{n}(s) \rangle \right\}. \quad (8)$$

where $a(s) = |n(s)|$ the local area of the surface at s . Thus, the \mathbb{L}^2 metric in the space of SRNFs is equivalent to a weighted sum of surface area change and surface bending. This property of SRNFs makes them very promising as a representation of surfaces for elastic shape analysis. 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} then there would be large gains in computational efficiency with respect to other metrics such as those defined on the space of thin shells [15].

The second important property of the SRNF map is that diffeomorphisms act on \mathcal{Q} by isometry, *i.e.* $\forall \gamma \in \Gamma, \|f_1 - f_2\| = \|f_1 \circ \gamma - f_2 \circ \gamma\|$. This will enable parameterization-invariant shape analysis and joint comparison, registration, and computation of geodesic, under the same Riemannian metric, between surfaces.

E. SRNF inversion

By performing all the computations required for correspondences, geodesics, and statistics, in \mathcal{Q} , which is Euclidean, and then mapping the results back to \mathcal{F} for visualization, the computational cost is reduced by an order of magnitude compared to methods that use, for example, path straightening under the pullback metric of the elastic metric. Unfortunately, there is no analytical expression for Q^{-1} for arbitrary points in \mathcal{Q} . Moreover, the injectivity and surjectivity of Q remain to be determined, meaning that for a given $q \in \mathcal{Q}$, there may be no $f \in \mathcal{C}_f$ such that $Q(f) = q$, and if such an f does exist, it may not be unique.

If one cannot invert the representation, one can always pull the \mathbb{L}^2 metric back to \mathcal{C}_f under Q and perform computations there, as in [23], but this is computationally expensive, and rather defeats the purpose of having an \mathbb{L}^2 metric in the first place. To solve this inversion problem, Xie et al. [24] showed that in the case of star-shaped surfaces, an analytical solution to the inversion problem exists. Laga et al. [25], on the other hand, developed 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 terms of the elastic metric, to q if it does not. This is achieved by formulating SRNF inversion as an optimization problem: find an element $f \in \mathcal{C}_f$ whose image $Q(f)$ is as close as possible to the given $q \in \mathcal{Q}$ under the \mathbb{L}^2 norm. This optimization problem is efficiently solved using a gradient descent approach [25]. To avoid undesirable solutions, since

gradient descent procedures converge to local minima if not initialized appropriately, Laga et al. [25] carefully engineered an orthonormal basis of \mathcal{C}_f , and used a spherical-wavelet based multiresolution and multiscale representation of the elements of \mathcal{C}_f . The basis can be generic, e.g. spherical harmonics, or domain-specific such as PCA basis in the presence of training datasets. With this approach, SRNF maps can be inverted with high accuracy, robustness to local minima, and low computational complexity.

Finally, if f_1 and f_2 are star-shaped surfaces, one can use the analytical solution of Xie et al. [24] as an initial guess for the numerical inversion procedure, thereby better initializing the reconstruction-by-optimization problem.

III. APPLICATIONS

In this section, we demonstrate a few applications of the tools discussed in this paper. Note that this section is not intended to provide a thorough performance evaluation but instead, it aims to show their potential and the range of shape analysis problems they can solve.

A. Elastic registration and geodesic deformation

The first example we consider is the joint registration and geodesic deformation of 3D objects that undergo complex elastic deformations. We emphasize that elastic (or non-rigid) deformation is composed of bending and stretching. Works based on (generalized) multi-dimensional scaling [31] and (heat-)diffusion geometry [32] only consider bending, which preserves intrinsic properties of surfaces. Here, we consider full elastic deformations.

Fig. 2 shows an example of joint elastic registration and geodesic computation between two human body shapes of the same subject but in significantly different poses. We first compare the quality of the registration when computed with a re-parameterization invariant elastic metric (Fig. 2-(b)) and when computed with other methods, e.g. functional maps of [33] (Fig. 2-(a)). As one can see, the latter method does not produce correct one-to-one correspondences compared to our approach, which formulates registration as a re-parameterization problem. Next, we compare geodesics computed using the \mathbb{L}^2 metric in \mathcal{C}_f vs. geodesics computed using elastic metrics by SRNF inversion. This is shown in Figs. 3-(a) and (b) where shapes undergo complex bending and stretching. Fig. 3-(c), on the other hand, shows a geodesic between two carpal bones. Fig. 1-(c) illustrates, for the same pair of 3D bones, the computed correspondences between a few landmarks.

B. Classification and clustering

The length of a geodesic, or deformation energy, between a pair of surfaces in the shape space is a measure of their dissimilarity that can be used for unsupervised as well as supervised classification of 3D shapes. To demonstrate this, we use the SHREC07 watertight 3D model benchmark [35], 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, triangulated meshes.

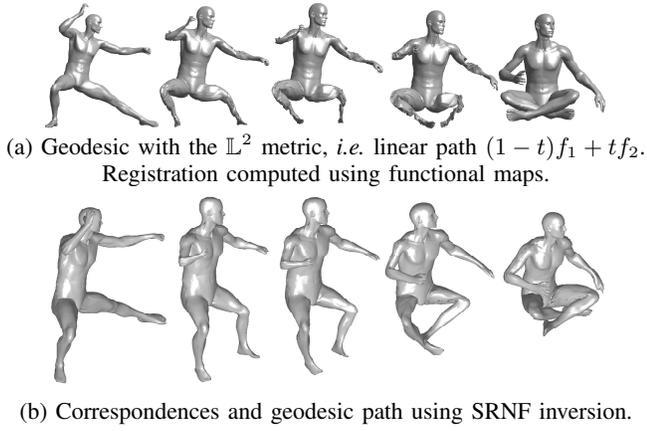


Fig. 2. Comparison between correspondences and geodesic computation using functional maps (a) and the SRNF inversion framework discussed in this paper (b).

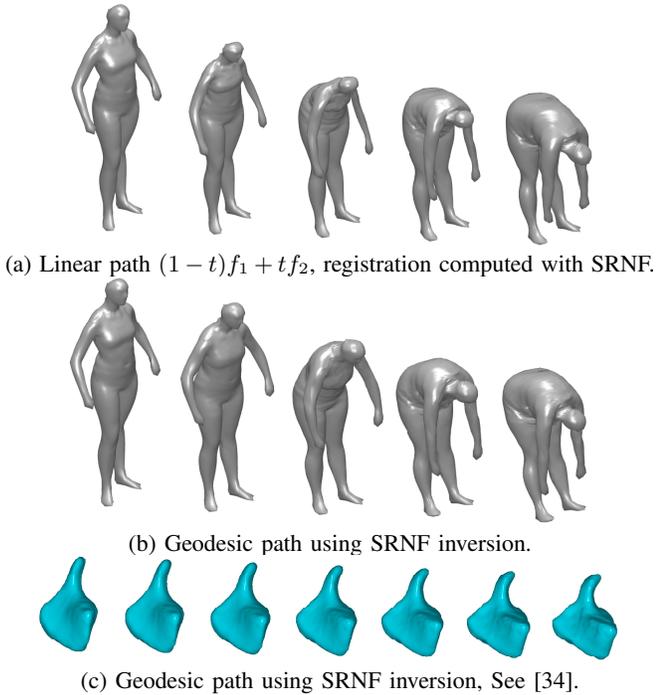


Fig. 3. Comparison between: (a) linear paths in \mathcal{F} ; (b) and (c) geodesics obtained by SRNF inversion of linear paths in \mathcal{Q} .

First, we compute their spherical parameterizations [36], and then normalize them for translation and scale. Next, we map the surfaces to the space of SRNFs, \mathcal{Q} , and compute their pairwise distance matrix in \mathcal{S}_q . That is, for a given pair of surfaces f_1 and f_2 and their respective SRNF maps q_1 and q_2 , their dissimilarity is given by

$$d(f_1, f_2) = d_S([f_1], [f_2]) = \min_{O, \gamma} \|q_1 - O(q_2 \circ \gamma)\|^2.$$

For comparison, we also compute the pairwise distances between the shapes using (1) the \mathbb{L}^2 distance in \mathcal{C}_f , *i.e.* without registration, and (2) the \mathbb{L}^2 metric in \mathcal{S}_f after elastic registration using the SRNF framework. Fig. 4 shows the multidimensional scaling (MDS) plots of these three pairwise

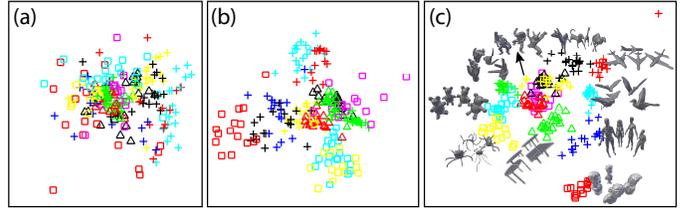


Fig. 4. MDS plots of the pairwise distances between every pair of models in the SHREC07 watertight database using different representations and metrics: (a) \mathbb{L}^2 metric in \mathcal{C}_f without registration; (b) \mathbb{L}^2 metric in \mathcal{S}_f with SRNF registration; (c) \mathbb{L}^2 metric in \mathcal{S}_q with SRNF registration. (Image best viewed in color.)

distance matrices. As one would expect, when using the \mathbb{L}^2 metric in \mathcal{C}_f , the 3D models are spread across the space (Fig. 4-(a)). Clusters start to emerge when using the \mathbb{L}^2 metric in \mathcal{S}_f with elastic registration performed using the SRNF framework (Fig. 4-(b)). When using the \mathbb{L}^2 metric in \mathcal{S}_q , which is equivalent to the partial elastic metric that quantifies bending and stretching (Fig. 4-(c)), we can clearly see that the 3D models are clustered by shape classes. This suggests that the SRNF framework is suitable for 3D shape classification and clustering.

C. Statistical 3D shape analysis and synthesis

Computing shape summaries on non-linear shape spaces using complex elastic metrics, or deformation energies, requires computing geodesics many times. At each iteration n geodesics must be computed, where n is the number of shapes in the collection to average. This is computationally very expensive as discussed in [25]. SRNF representations offer an elegant and a computationally efficient alternative solution. Instead of using the expensive pullback metric, one can simplify the process by first mapping all the surfaces, $\{f_1, \dots, f_n\}$, to the SRNF space, \mathcal{Q} , resulting in $\{q_1, \dots, q_n\}$. Then, the mean shape in SRNF space, denoted by \bar{q} , is computed by iterating between co-registering all q_i s to \bar{q} , and subsequently updating \bar{q} . This is computationally very efficient since \mathcal{Q} is a Euclidean space. Finally, at the last step, the mean shape $\bar{f} \in [\bar{f}]$ is computed by SRNF inversion, *i.e.* finding a shape \bar{f} such that $Q(\bar{f}) = \bar{q}$. The entire procedure is summarized in Algorithm 1.

Algorithm 1: Sample Karcher mean by SRNF inversion.

Input: Set of surfaces $\{f_1, \dots, f_n\} \in \mathcal{C}_f$ and their SRNFs $\{q_1, \dots, q_n\} \in \mathcal{Q}$.

Output: Karcher mean surface \bar{f} .

- 1: Let $\bar{q} = Q(\bar{f})$ with \bar{f} set to f_1 as the initial estimate of the Karcher mean. Set $j = 0$.
 - 2: For each $i = 1, \dots, n$, register q_i to \bar{q} resulting in $q_i^* = O_i^*(q_i, \gamma_i^*)$, where O_i^* and γ_i^* are the optimal rotation and re-parameterization, respectively.
 - 3: Update the average $\bar{q} = \frac{1}{n} \sum_{i=1}^n q_i$.
 - 4: If change in $\|\bar{q}\|$ is small, stop. Otherwise go to Step 2.
 - 5: Find \bar{f} by SRNF inversion such that $Q(\bar{f}) = \bar{q}$.
-

Next, let \bar{q} denote the SRNF of the average shape. With a slight abuse of notation, let q_i , $i = 1, \dots, n$ denote the SRNFs of the surfaces in the sample optimally registered to this average. Then, since the SRNF space is Euclidean, the modes of variation can be computed in a standard way, *i.e.* using singular value decomposition (SVD) of the covariance matrix of the SRNF maps. Let u_q^k denote the k -th principal direction. Then, one can explore the variability in this direction around the mean using $\bar{q} + \lambda u_q^k$, where $\lambda \in \mathbb{R}$. To visualize the principal directions of variation in \mathcal{F} , we only need to compute, by SRNF inversion, f_λ^k such that $Q(f_\lambda^k) = \bar{q} + \lambda u_q^k$.

Fig. 5 shows the mean shape and the first four modes of variation produced with this procedure on a collection of human body shapes. An important by-product of this procedure is the simultaneous co-registration of multiple shapes. In fact, in the process of computing the mean shape, all the surfaces in the collection are registered to the final mean shape and subsequently registered to each other.

Finally, we consider the problem of synthesizing arbitrary 3D models that are similar to, but not exactly the same as, a given collection of 3D models. We consider a collection of 398 human shapes [37] composed of multiple subjects in different poses. The first subject has 35 different poses including a neutral pose and a few other poses. We first compute the SRNF representations of these surfaces, perform statistical analysis in the SRNF space \mathcal{Q} , using standard linear statistics such as Principal Component Analysis, and finally map the results back to the surface space, \mathcal{F} , using the SRNF inversion algorithm. Formally, if \bar{q} is the mean and $V_i, i = 1, \dots, k$ are the leading eigenvectors in the space of SRNFs, then any arbitrary SRNF map q can be written as: $q = \bar{q} + \sum_{i=1}^k \lambda_i V_i$, where λ_i are real coefficients. An arbitrary shape can then be synthesized by finding a surface f such that its SRNF is q . Fig. 6 shows 10 randomly synthesized human body shapes. Observe that the statistical model can generate shapes of arbitrary subjects in arbitrary poses.

D. Analysis of planar shapes of arbitrary topology

Finally, we show how the proposed framework can be used for computing geodesics between planar shapes of arbitrary topology. Most previous Riemannian formulations of 2D shape analysis are restricted to curves that can be parameterized with a single parameter domain [38], [39], [40]. However, 2D shapes may contain multiple connected components and many internal details that cannot be captured by such parameterizations. We propose to represent planar objects as level sets of their Euclidean distance functions. We consider each such function along with its smoothed gradient as a 2D surface $f : [0, 1]^2 \rightarrow \mathbb{R}^3$. We thus formulate the problem of analyzing curves as a problem of elastic shape analysis of surfaces, and adapt the Riemannian framework described above to the problem at hand. Fig. 7 shows two examples of geodesics between pairs of planar shapes of arbitrary topology computed using this approach. For further results and discussion, we refer the reader to [41].

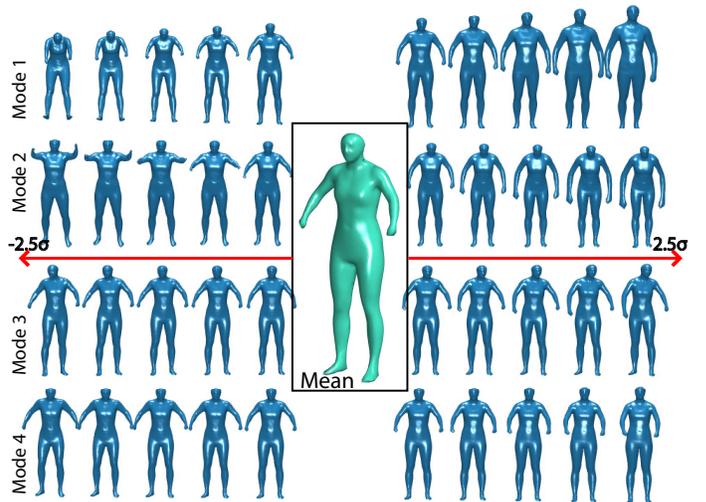


Fig. 5. Mean (central shape) and first five leading modes of variations, computed using SRNF inversion, of a collection of 3D human body shapes in different poses and belonging to different subjects.

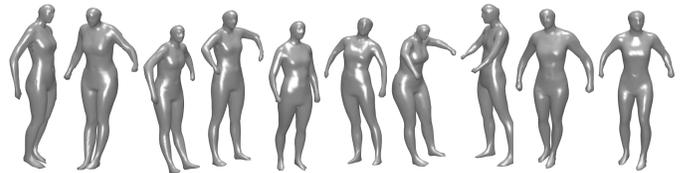


Fig. 6. Ten arbitrary 3D human body shapes automatically synthesized by sampling from a Gaussian distribution fitted to a collection of human body shapes in SRNF space.

IV. CONCLUSION

In this article, we have summarized and discussed a framework based on square-root normal fields (SRNF) for the statistical analysis of the shape of 3D objects that undergo full elastic deformations. The framework offers several theoretical and computational benefits compared to other methods. First, the \mathbb{L}^2 metric in the space of SRNFs is equivalent to the partial elastic metric. Second, the diffeomorphism group acts by isometries on the space of SRNFs. These two properties enable us to simultaneously register surfaces and compute geodesics between surfaces that undergo large bending and stretching. We demonstrated several applications of this framework. Possible directions for future work is the generalization of this framework to surfaces of arbitrary topology and to 3D shapes that undergo structural deformations.

REFERENCES

- [1] U. Grenander and M. I. Miller, “Computational anatomy: an emerging discipline,” *Q. Appl. Math.*, vol. LVI, no. 4, pp. 617–694, Dec. 1998.
- [2] J. Fishbaugh, M. Prastawa, S. Durrleman, J. Piven, and G. Gerig, “Analysis of longitudinal shape variability via subject specific growth modeling,” in *MICCAI*, 2012, pp. 731–738.
- [3] D. G. Kendall, “The diffusion of shape,” *Advances in Applied Probability*, vol. 9, no. 3, pp. 428–430, 1977.
- [4] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or, “A survey on shape correspondence,” *Comp. Graph. Forum*, vol. 30, no. 6, pp. 1681–1707, 2011.

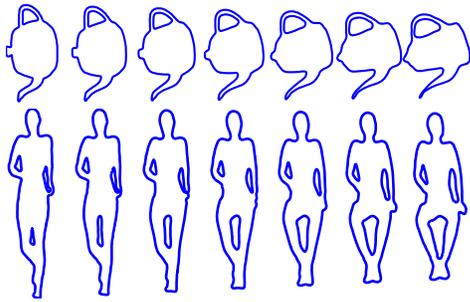


Fig. 7. Geodesics between planar shapes of arbitrary topology.

[5] H. Laga, H. Takahashi, and M. Nakajima, "Spherical wavelet descriptors for content-based 3D model retrieval," in *IEEE International Conference on Shape Modeling and Applications 2006*, 2006, pp. 15–.

[6] H. Tabia, H. Laga, D. Picard, and P.-H. Gosselin, "Covariance descriptors for 3d shape matching and retrieval," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 4185–4192.

[7] H. Tabia and H. Laga, "Covariance-based descriptors for efficient 3d shape matching, retrieval and classification," *IEEE Transactions on Multimedia*, vol. 17, no. 9, pp. 1591–1603, 2015.

[8] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*. John Wiley & Sons, 1998.

[9] B. Allen, B. Curless, and Z. Popović, "The space of human body shapes: reconstruction and parameterization from range scans," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 587–594, July 2003.

[10] S. Bouix, J. Pruessner, D. Collins, and K. Siddiqi, "Hippocampal shape analysis using medial surfaces," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2001, vol. 2208, pp. 33–40.

[11] K. Gorcowski, M. Styner, J. Y. Jeong, J. Marron, J. Piven, H. C. Hazlett, S. M. Pizer, and G. Gerig, "Multi-object analysis of volume, pose, and shape using statistical discrimination," *IEEE trans. on PAMI*, vol. 32, no. 4, pp. 652–661, 2010.

[12] S. Osher and N. Paragios, *Geometric level set methods in imaging, vision, and graphics*. Springer, 2003.

[13] R. H. Davies, C. J. Twining, T. F. Cootes, and C. J. Taylor, "Building 3-d statistical shape models by direct optimization," *IEEE Trans. on Medical Imaging*, vol. 29, no. 4, pp. 961–981, 2010.

[14] M. Kilian, N. J. Mitra, and H. Pottmann, "Geometric modeling in shape space," *ACM Trans. on Graphics*, vol. 26, no. 3, July 2007.

[15] B. Heeren, M. Rumpf, M. Wardetzky, and B. Wirth, "Time-discrete geodesics in the space of shells," *CGF*, vol. 31, no. 5, pp. 1755–1764, Aug. 2012.

[16] C. Zhang, B. Heeren, M. Rumpf, and W. A. Smith, "Shell pca: statistical shape modelling in shell space," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1671–1679.

[17] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. Van Kaick, and A. Tagliasacchi, "Deformation-driven shape correspondence," in *Computer Graphics Forum*, vol. 27, no. 5. Wiley Online Library, 2008, pp. 1431–1439.

[18] T. Windheuser, U. Schlickewei, F. R. Schmidt, and D. Cremers, "Geometrically consistent elastic matching of 3d shapes: A linear programming solution," in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 2134–2141.

[19] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of closed surfaces for 3-d shape description," *CVIU*, vol. 61, no. 2, pp. 154–170, 1995.

[20] M. Styner, I. Oguz, S. Xu, C. Brechbühler, D. Pantazis, J. J. Levitt, M. E. Shenton, and G. Gerig, "Framework for the statistical shape analysis of brain structures using spharm-pdm," *The insight journal*, no. 1071, p. 242, 2006.

[21] D. A. Hirshberg, M. Loper, E. Rachlin, and M. J. Black, "Coregistration: Simultaneous alignment and modeling of articulated 3D shape," in *European Conference on Computer Vision*. Springer, 2012, pp. 242–255.

[22] I. H. Jermyn, S. Kurtek, E. Klassen, and A. Srivastava, "Elastic shape matching of parameterized surfaces using square root normal fields," in *ECCV*, 2012, pp. 804–817.

[23] Q. Xie, S. Kurtek, H. Le, and A. Srivastava, "Parallel transport of deformations in shape space of elastic surfaces," in *Computer Vision, IEEE International Conference on*, December 2013.

[24] Q. Xie, I. Jermyn, S. Kurtek, and A. Srivastava, "Numerical inversion of smfs for efficient elastic shape analysis of star-shaped objects," in *European conference on computer vision*. Springer International Publishing, 2014, pp. 485–499.

[25] H. Laga, Q. Xie, I. H. Jermyn, and A. Srivastava, "Numerical inversion of smf maps for elastic shape analysis of genus-zero surfaces," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.

[26] I. Dryden and K. Mardia, *Statistical Shape Analysis*, 1998.

[27] A. Bhattacharya and R. Bhattacharya, "Statistics on Riemannian manifolds: asymptotic distribution and curvature," *Proceedings of the American Mathematical Society*, vol. 136, no. 8, pp. 2959–2967, 2008.

[28] C. Small, "The statistical theory of shape," 1996.

[29] S. H. Joshi, Q. Xie, S. Kurtek, A. Srivastava, and H. Laga, "Surface shape morphometry for hippocampal modeling in alzheimer's disease," in *DICTA*, 2016, p. (to appear).

[30] P. G. Ciarlet, "An introduction to differential geometry with applications to elasticity," *Journal of Elasticity*, vol. 78, no. 1-3, pp. 1–215, 2005.

[31] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, no. 5, pp. 1168–1172, 2006.

[32] M. M. Bronstein and I. Kokkinos, "Scale-invariant heat kernel signatures for non-rigid shape recognition," in *International Conference on Computer Vision and Pattern Recognition*. IEEE, 2010, pp. 1704–1711.

[33] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas, "Functional maps: a flexible representation of maps between shapes," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 30:1–30:11, July 2012.

[34] I. Banerjee, H. Laga, G. Patané, S. Kurtek, A. Srivastava, and M. Spagnuolo, "Generation of 3d canonical anatomical models: an experience on carpal bones," in *International Conference on Image Analysis and Processing*. Springer International Publishing, 2015, pp. 167–174.

[35] D. Giorgi, S. Biasotti, and L. Paraboschi, "Shape retrieval contest 2007: Watertight models track," *SHREC competition*, vol. 8, 2007.

[36] S. Kurtek, A. Srivastava, E. Klassen, and H. Laga, "Landmark-guided elastic shape analysis of spherically-parameterized surfaces," *Comp. Graph. Forum*, vol. 32, no. 2pt4, pp. 429–438, 2013.

[37] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel, "A statistical model of human pose and body shape," in *CGF*, vol. 28, no. 2, 2009, pp. 337–346.

[38] A. Srivastava, E. Klassen, S. Joshi, and I. Jermyn, "Shape analysis of elastic curves in euclidean spaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 7, pp. 1415–1428, 2011.

[39] H. Laga, S. Kurtek, A. Srivastava, M. Golzarian, and S. J. Miklavcic, "A riemannian elastic metric for shape-based plant leaf classification," in *Digital Image Computing Techniques and Applications (DICTA)*, 2012 *International Conference on*. IEEE, 2012, pp. 1–7.

[40] H. Laga, S. Kurtek, A. Srivastava, and S. J. Miklavcic, "Landmark-free statistical analysis of the shape of plant leaves," *Journal of theoretical biology*, vol. 363, pp. 41–52, 2014.

[41] S. Kurtek, H. Laga, and Q. Xie, "Elastic shape analysis of boundaries of planar objects with multiple components and arbitrary topologies," in *Asian Conference on Computer Vision*. Springer International Publishing, 2014, pp. 424–439.