



HAL
open science

Gauge Invariant Framework for Shape Analysis of Surfaces

Alice Barbara Tumpach, Hassen Drira, Mohamed Daoudi, Anuj Srivastava

► **To cite this version:**

Alice Barbara Tumpach, Hassen Drira, Mohamed Daoudi, Anuj Srivastava. Gauge Invariant Framework for Shape Analysis of Surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016, 38 (1), pp.46 - 59. hal-01142780

HAL Id: hal-01142780

<https://hal.science/hal-01142780>

Submitted on 16 Apr 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Gauge Invariant Framework for Shape Analysis of Surfaces

Alice Barbara Tumpach, Hassen Drira, Mohamed Daoudi, *Senior, IEEE*,
and Anuj Srivastava, *Senior, IEEE*

Abstract—This paper describes a novel framework for computing geodesic paths in shape spaces of spherical surfaces under an elastic Riemannian metric. The novelty lies in defining this Riemannian metric directly on the quotient (shape) space, rather than inheriting it from pre-shape space, and using it to formulate a path energy that measures only the normal components of velocities along the path. In other words, this paper defines and solves for geodesics directly on the shape space and avoids complications resulting from the quotient operation. This comprehensive framework is invariant to arbitrary parameterizations of surfaces along paths, a phenomenon termed as gauge invariance. Additionally, this paper makes the link between different elastic metrics used in the computer science literature on one hand, and the mathematical literature on the other hand, and provides a geometrical interpretation of the terms involved. Examples using real and simulated 3D objects are provided to help illustrate the main ideas.

Index Terms—3D surfaces, Riemannian metric, geodesics.

1 INTRODUCTION

In this paper we seek a framework for analysing shapes of a certain class of 3D objects. Although the general goal in shape analysis is to develop tools for full statistical analysis – statistical averaging, finding principal modes of variations in a population, and shape classification, we restrict to more basic goals of quantifying shape differences and generating deformations. While there have been many efforts in shape analysis of 3D objects, the problem is far from solved and the current solutions face many technical and practical issues. For instance, many general techniques for shape analysis rely on quantifying shape differences by spatially matching geometric features across objects. Therefore, it becomes important to establish a correspondence of parts between objects, i.e. which part in one object corresponds to which part in the other? This was an important bottleneck in a majority of previous efforts on 3D shape analysis where the correspondence (or registration) of objects was either presumed or solved as an independent pre-processing step. More recently, there has been progress in establishing frameworks that formulate the registration and comparison problems jointly. These newer frameworks, using

techniques from differential geometry, focus on shape analysis of parameterized surfaces and treat the problem of shape comparison as the problem of computing geodesic paths in shape spaces under a chosen metric. Here shapes are compared using a Riemannian metric on a *pre-shape space* \mathcal{F} consisting of embeddings or immersions of a model manifold (like the sphere, or the disc) into the 3D Euclidean space \mathbb{R}^3 . Two embeddings correspond to the same shape in \mathbb{R}^3 if and only if they differ by an element of a shape-preserving transformation group, such as rigid motion, scaling, and reparameterization. The shape space is therefore the quotient space of the pre-shape space by these shape-preserving groups. If the Riemannian metric on the pre-shape space is preserved by the action of the shape-preserving group then it induces a Riemannian metric on the quotient space. The construction of geodesics in shape space provide optimal deformations between surfaces and is a very important tool in *statistical* analysis of shapes. Interestingly, the problem of registration is handled using parameterizations of surfaces such that the points denoting the same parameter values on two objects are considered registered.

While these geometric ideas are powerful and comprehensive, there are two important issues that one needs to deal with: (1) the choice of Riemannian metric to define geodesics, geodesic lengths, and the eventual shape metric, and (2) the task of computing geodesic paths between arbitrary shapes. In terms of the first issue, the choice of a metric, an important requirement is that the metric should be invariant to action of the reparameterization group, to enable a well-defined distance on the eventual quotient space or the

- A.B. Tumpach is with Laboratoire Paul Painlevé (U.M.R. CNRS 8524) University of Lille1, France. E-mail: Barbara.Tumpach@math.univ-lille1.fr
- H. Drira is with Institut Mines-Telecom/Telecom Lille CRISAL (UMR CNRS 9189). E-mail: hassen.drira@telecom-lille.fr
- M. Daoudi is with Institut Mines-Telecom/Telecom Lille CRISAL (UMR CNRS 9189). E-mail: mohamed.daoudi@telecom-lille.fr
- A. Srivastava is with the Department of Statistics, Florida State University, Tallahassee, USA. E-mail: anuj@fsu.edu

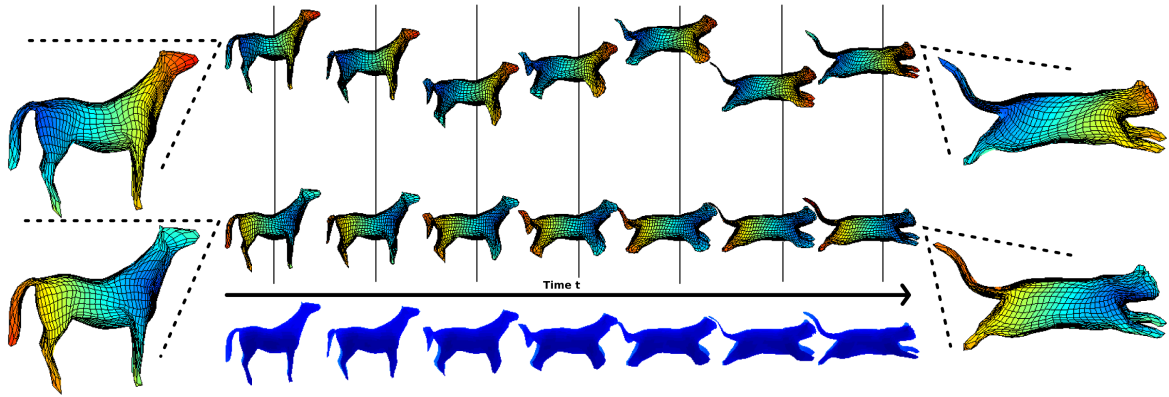


Fig. 1. Two paths in \mathcal{F} with the same sequence of shapes but with different reparameterizations between the corresponding shapes.

shape space of surfaces. There is a related requirement for the shape analysis to be invariant to parameterizations of objects since parameterizations are only artificial impositions designed to help navigate along objects. The physical intuition we have is that shape tools, such as the deformation (path or geodesic) from one shape to another, are physical processes that are independent of the way surfaces may be parameterized. These dual requirements rule out the use of commonly-used quantities such as the \mathbb{L}^2 norm on the space \mathcal{F} directly. In terms of the second issue, the lack of standard metrics makes it complicated to compute geodesic paths even when the underlying manifold is a vector space, and one needs numerical algorithms for approximating geodesic paths. Next, we present a summary of the past work on these two issues and outline motivations for the current paper.

1.1 Motivation and Past Work

Our goal is to develop tools for analyzing shapes of two-dimensional surfaces with certain local constraints (smoothness, no-holes, etc). The main difficulty in comparing shapes of such surfaces is that there is no preferred parameterization that can be used for registering and comparing features across surfaces. Since the shape of a surface is invariant to its parameterization, one would like an approach that yields the same result irrespective of the parameterization.

Furthermore, we are not only interested in the comparison and matching of two shapes, but also in the deformation processes that may transform one shape into another, i.e. metamorphosis. To be physically meaningful, the evolution from one shape to another should be independent of the way surfaces may be parameterized. Our approach to shape analysis presented in this paper was therefore initiated by the following question: *What is the natural framework where one can measure deformations of shapes independently of the way shapes are parameterized?* As a motivating example, the sequence of shapes displayed in Fig. 1

(bottom) denotes a path where a horse is transformed into a jumping cat. During the transformation process, only the change of shape, drawn in the bottom line as a sequence of blue surfaces, is relevant to us. How the surfaces may be parameterized during the metamorphosis has no importance in our context. To emphasize this idea, two paths of parameterized surfaces corresponding to the same transformation process are displayed in the top two rows. We would like a framework where the physical quantities measured on the path of shapes, such as its length or its energy, are independent of the parameterizations of surfaces along the transformation process. In particular, in Fig. 1, the two paths of parameterized surfaces corresponding to the same transformation process should have the same length. Note that the surfaces along the second path are obtained by applying to the surfaces along the first path a *different reparameterization at each time step*.

Let us emphasize that we are not only interested in how *far* the horse and the jumping cat are from each other, in other words in a quantity like a distance measuring the minimal cost needed to deform the horse into a cat. But, given a metamorphosis between these two shapes, we are also interested in measuring its *length* on one hand, and its *energy* on the other hand, independently of the parameterizations of the transformation process that may have been used to create this metamorphosis. Recall that the length of a path is the integral of the norm velocity function with respect to time and has the dimension of a distance. The energy is the integral of the square of the norm velocity function with respect to time, hence has the dimension of the square of a distance divided by time.

Let us now summarize past work on related subjects. The initial set of papers developed algorithms for geodesic deformations between surfaces while using the given registration of points. They compute geodesics between shapes, under isometric deformations, while assuming the registration (or parameterization) as given. Windheuser et al. [1] proposed to

find a geometrically consistent matching of 3D shapes which minimizes an elastic deformation energy but use a linear interpolation between registered pairs of points in \mathbb{R}^3 to compute geodesic paths. Another paper by Kilian et al. [2] represents parameterized surfaces by discrete triangulated meshes, assumes a Riemannian metric on the space of such meshes, and computes geodesic paths between given meshes. The main limitation here is that it assumes the correspondence between points across meshes. That is, we need to know beforehand which point on one mesh matches with which point on the second mesh. The same limitation holds for the paper by Heeren et al. [3] also. In contrast, we would like to remove the reparameterization variability so that different surfaces with the same shape but different parameterizations have zero distance between them.

Motivated by progress in shape analysis of curves [4], [5], Kurtek et al. [6], [7] introduced a new representation, termed a q -map of surfaces such that the \mathbb{L}^2 distance in this representation space is invariant to simultaneous reparameterizations of surfaces. For convenience of the reader, we recall the definition of the q -map but we will not use it in the present paper. Let $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$ denote a smooth parameterized surface and \mathcal{F} be the set of such surfaces. Then, this q -map is given by $f \mapsto q$ where $q(s) = \sqrt{r(s)}f(s)$ and $r(s)$ is the area multiplication factor of f at $s \in \mathbb{S}^2$. They defined a Riemannian metric on the space of parameterized surfaces by *pulling back* \mathbb{L}^2 metric under the q -map, and used a *path-straightening* algorithm to compute geodesic paths between given surfaces in a pre-shape space. This path-straightening is an iterative algorithm that updates an arbitrary initial path using the gradient of the energy function mentioned above, until the path converges to a geodesic. The energy gradient is approximated numerically using an (approximate) finite basis for \mathcal{F} . To remove the effects of original parameterizations, and to obtain geodesics in the shape space, they solve for an optimal reparameterization of one of the surfaces, under the same energy. There are several other papers, including [8], that focus exclusively on the task of finding optimal correspondence between 3D objects, either using physically-motivated energies or Riemannian metrics. Due to the use of gradient-based searches, these methods and previously mentioned papers do not guarantee a global solution, either for geodesics or for registration. In path-straightening, however, it can be shown that a path that is a local minimum of the path energy is a geodesic path, albeit not the shortest geodesic. To our knowledge, very few methods guarantee a globally-optimal solution to the problem of finding geodesics in shapes spaces of surfaces. Although paper [6] was the first to provide a geometric framework for joint registration-comparison problem, the Riemannian metric used there has a limitation that it was not translation invariant.

To handle the translation issue mentioned above, Jermyn et al. [9] introduced a comprehensive Riemannian metric that has several improvements, including the fact that it was translation invariant and allows some physical interpretations in its use. This metric, given later in Eqn. (7), has terms that can be interpreted as measurements of bending, stretching, and changes in local curvatures of surfaces. (We elaborate on this topic later in Section 3.4.) It has been termed an *elastic metric* because it is invariant to reparameterizations and the physical interpretations associated with it. Although [9] introduced this metric, it did not use the full metric to compute geodesic paths. Instead, it defined a new map, termed the square-root normal field, given by $q(s) = \sqrt{r(s)}n_f(s)$ where $n_f(s)$ denotes the unit normal to the surface at the point $s \in \mathbb{S}^2$. The square-root normal field has the property that the *last two terms* of the elastic metric transform to the \mathbb{L}^2 metric under the map $f \mapsto q$, for some weighting of last two terms in the metric. The first term of the metric is discarded in this analysis. The transformation to \mathbb{L}^2 metric is useful since one can apply some common tools from Hilbert space analysis to this problem, including the optimization over the reparameterization group for optimal registration, but this mapping $f \mapsto q$ is not onto and, hence, not invertible. The optimization step is challenging because the reparameterization group is an infinite-dimensional Fréchet Lie group, and the exponential map is not a local diffeomorphism. Since the first term of the elastic metric introduced in [9] is not used by Jermyn et al., it can result in zero shape distance between two surfaces that actually have different shapes. For example, a thin-tall cylinder and a fat-short cylinder, with same surface areas and unit normals, will have zero shape difference under this framework.

Another line of work in shape analysis comes from Michor et al. [10], Bauer et al. [11], [12], [13] and Fuchs et al. [14] (see also [15] for an overview of a lot of mathematical results in this area). Different types of metrics have been studied : Sobolev metrics in [13], curvature weighted metrics in [11], almost local metrics in [12], metrics measuring the deformations of the interiors of shapes in [14]. Let us mention that the first two terms of the metric we use in the present paper fit in the general study laid out in [13], and are related to the metrics studied in [16], [17], [18] (in a sense that we will make clear in section 3.4). In this set of papers, the idea is to replace the problem of solving the geodesic equation on shape space by the equivalent problem of solving the equation for horizontal geodesics in the pre-shape space. A geodesic in pre-shape space is horizontal if it is orthogonal to the orbits of the reparameterization group. One task in this strategy is therefore to compute the horizontal space on which the quotient map is an isometry, or equivalently solve a minimization problem for the infinitesimal energy. Depending on the Riemannian

metric on the pre-shape space, this task may be computationally trivial or extremely difficult to implement (for metrics used in [11] and [12] it is just the space of normal vector fields, but for metrics used in [13] and [14] it involves the inversion of a pseudo-differential operator). Another main contribution of these authors is to give sufficient conditions under which the Riemannian metric induced on shape space separates points, i.e. gives a non-zero geodesic distance between pairs of different shapes (a condition that is necessary to make shape comparison). It is worth noting that, in this infinite-dimensional context, vanishing geodesic distance is a common phenomenon (as was first highlighted in [10]). For the metric we use, non-vanishing geodesic distance is guaranteed by the non-vanishing geodesic distance on the space of Riemannian metrics proved in [19] (at least on pairs of shapes inducing different pull-back metrics on the sphere, which is what we are interested in practice).

To summarize, the past approaches involving Riemannian geometry have tended to perform shape analysis in two steps. First, they select a representation space, or a pre-shape space, for objects of interest – curves [4], [5], [20] and surfaces [7], [9], [11], [12], [13] – and impose a Riemannian structure on it ensuring that the actions of shape-preserving groups are by isometries. Next, they inherit this metric to the quotient space of the pre-shape space modulo the requisite groups, called the shape space, and seek geodesics between objects in this shape space. The task of inheriting Riemannian metrics to quotient spaces is complicated because reparameterization groups are Fréchet Lie groups and the process of inheriting a metric requires closed orbits, as can be seen in [21], [4], [20], [13], etc. Even though endowing shape space with a Riemannian metric (with positive distance function) seems to be a good approach, inducing this metric by a Riemannian metric on pre-shape space leads to difficulties that one would like to avoid (recall that we are only interested in shapes and not in the way they are parameterized). We will pursue a different strategy where the Riemannian metric is directly imposed on the quotient space, thus avoiding the need to satisfy conditions for inheriting metrics from the pre-shape space or computing an abstract horizontal space. Motivated by an easy implementation of the metrics, we take the point of view where the space of interest is the space of normal vector fields (in contrast with the horizontal space of a Riemannian submersion). Let us emphasise that there is no restriction in doing so : any Riemannian metric on shape space can be expressed as a metric defined on normal vector fields.

1.2 Goals and Contributions

Now we present the goals and contributions of this paper, and start by revisiting the question: *What should be a good Riemannian metric on shape space ?* A good

Riemannian metric on shape space should be such that : (1) it induces a positive distance function on shape space, i.e. the infimum of the lengths of paths connecting two different shapes should be non-zero ; (2) the distance between two shapes should be independent of the way the two shapes are parameterized ; and, (3) the length of a path of shapes should be independent of the way shapes along the path are parameterized. The last point should be thought of as the natural generalization of the fact that, on a finite-dimensional Riemannian manifold, the length of a curve is independent of the way the curve is parameterized. It should be true for any path (not only for geodesics), and is called *gauge invariance*. Indeed the use of parameterized surfaces in order to measure the deformation of a shape can be compared to the use of a gauge. Let us comment on Fig. 1 in order to illustrate this idea. Each column depicts an orbit under the reparameterization group for the corresponding surface, the surfaces in a given orbit correspond to the same shape but with different parameterizations. A path of shapes can be lifted in many ways to a path of parameterized surfaces. In Fig. 1 two lifts of the bottom line path are depicted. The first path connects parameterized surfaces with different "heights" in the fibers. This is made to emphasize that the variations of the "height" (i.e. of the parameterization) in the fibers should not influence the value of the length of the path of shapes.

The main contributions of this paper are following:

- The proposed method achieves gauge invariance, i.e. the lengths of paths (geodesics or otherwise) measured under this metric are invariant to arbitrary reparameterizations of shapes along these paths (in particular, the two paths in Fig.1 have the same length).
- It uses an elastic metric that accounts for any deformation of patches to define and compute geodesic paths between given objects in the shape space, and it presents a geometric interpretation of the different terms involved in this metric.
- By defining a metric directly in the shape space, it avoids the optimization step over the reparameterization group and difficult mathematical issues arising from inheriting a metric from pre-shape space.

Note that the third point leads to more efficient Algorithms in cases where one only needs a shape geodesic and not the optimal registration between surfaces. It provides the same geodesic path despite arbitrary initial parameterizations (or registrations) of given surfaces, and saves the computational cost of finding a registration. This fact is also a source of limitation in the situation where one needs a registration. If one wants to use geodesic lengths for comparing shapes, then a registration is not needed. However, if one wants to study statistical summaries of deformation

fields, then a registration will be needed.

The rest of this paper is organized as follows. Section 2 describes the mathematical representation of embedded surfaces and establishes mathematical setup. Section 3 is devoted to the description of gauge invariance and to the definition of the Riemannian metric involved in this paper. The geodesic computation is described in Section 4 and Section 5 presents the experimental results.

2 MATHEMATICAL SETUP

2.1 Notation

We will represent a shape S with an embedding $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$ such that the image $f(\mathbb{S}^2)$ is S . The function f is also called a *parameterization* of the surface S .

We will use local coordinates (u, v) on the sphere. For the theoretical framework, any coordinates on the sphere are suitable, but in the application we use spherical coordinates : u stands for the polar angle and ranges from 0 to π , and v denotes azimuthal angle and ranges from 0 to 2π .

Recall that a map $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$ is an *embedding* when: for any point $(u, v) \in \mathbb{S}^2$, (1) f is smooth, in particular the derivatives f_u and f_v of f with respect to u and v are well-defined, (2) f is an immersion, i.e. the cross product $f_u \times f_v$ never vanishes and allows us to define the normal (resp. tangent) space to the surface $f(\mathbb{S}^2)$ at a point $f(u, v)$ as the subspace of \mathbb{R}^3 which is generated by (resp. orthogonal to) $f_u \times f_v$, and (3) f is an homeomorphism onto its image, i.e. points on $f(\mathbb{S}^2)$ that look close in \mathbb{R}^3 are images of close points in \mathbb{S}^2 . If f is an embedding, then the surface $f(\mathbb{S}^2)$ is naturally oriented by the frame $\{f_u, f_v\}$, or equivalently by the normal vector field $f_u \times f_v$.

We define the space of all such surfaces as

$$\mathcal{F} := \{f : \mathbb{S}^2 \rightarrow \mathbb{R}^3, f \text{ is an embedding}\}.$$

It is often called the *pre-shape space* since objects with same shape but different orientations or parameterizations may correspond to different points in \mathcal{F} . The set \mathcal{F} is itself a manifold, as an open subset of the linear space $\mathcal{C}^\infty(\mathbb{S}^2, \mathbb{R}^3)$ of smooth functions from \mathbb{S}^2 to \mathbb{R}^3 (see Theorem 3.1 in [15] and the references therein). The tangent space to \mathcal{F} at f , denoted by $T_f\mathcal{F}$, is therefore just $\mathcal{C}^\infty(\mathbb{S}^2, \mathbb{R}^3)$.

The shape-preserving transformations of 3D object can be expressed as group actions on \mathcal{F} : The group \mathbb{R}^+ with multiplication operation acts on \mathcal{F} by *scaling* : $(\beta, f) \mapsto \beta f$, for $\beta \in \mathbb{R}^+$ and $f \in \mathcal{F}$. The group \mathbb{R}^3 with addition as group operation acts on \mathcal{F} , by *translations* : $(v, f) \mapsto f + v$, for $v \in \mathbb{R}^3$ and $f \in \mathcal{F}$. The group $\text{SO}(3)$ with matrix multiplication as group operation acts on \mathcal{F} , by *rotations* : $(O, f) \mapsto Of$, for $O \in \text{SO}(3)$ and $f \in \mathcal{F}$. Finally, the group $\Gamma := \text{Diff}^+(\mathbb{S}^2)$ consisting of diffeomorphisms which preserve the orientation of \mathbb{S}^2 acts also on \mathcal{F} , by *reparameterization* : $(\gamma, f) \mapsto f \circ \gamma^{-1}$, for $\gamma \in \text{Diff}^+(\mathbb{S}^2)$ and $f \in \mathcal{F}$. The use of γ^{-1} , instead

of γ , ensures that the action is from left and, since the action of $\text{SO}(3)$ is also from left, one can form a joint action of $G := \text{Diff}^+(\mathbb{S}^2) \times \text{SO}(3) \rtimes \mathbb{R}^3$ on \mathcal{F} . In this paper, the translation group is taken care of by using translation-independant metric (the elastic metric) and, when needed, the scaling is taken care of by rescaling the surfaces to have unit surface area. Therefore, in the following we will focus only on the reparameterization group Γ and on the rotation group $\text{SO}(3)$.

2.2 Shape Space as quotient space

Since we are only interested in shapes of surfaces, we would like to identify surfaces that can be related through a shape-preserving transformation. This is accomplished using the notion of group action and orbits under those group actions.

Given a group G acting on \mathcal{F} , the elements in \mathcal{F} obtained by following a fix parameterized surface $f \in \mathcal{F}$ when acted on by all elements of G is called the G -orbit of f or the *equivalence class* of f under the action of G , and will be denoted by $[f]$. In particular, when G is the reparameterization group, the orbit of $f \in \mathcal{F}$ is characterized by the surface $f(\mathbb{S}^2) = S$, i.e. the elements in $[f] = \{f \circ \gamma^{-1} \text{ for } \gamma \in \Gamma\}$ are all possible parameterizations of S . For instance in Fig. 1, the first column contains some parameterized horses that are elements of the same orbit. The set of orbits of \mathcal{F} under a group G is called the *quotient space* and will be denoted by \mathcal{F}/G . The quotient space of interest in this paper is called *shape space* and is defined as follows.

Definition 1: The *shape space* \mathcal{S} is the set of oriented surfaces in \mathbb{R}^3 , which are diffeomorphic to \mathbb{S}^2 , modulo translation and rotation. It is isomorphic to the quotient space of the pre-shape space \mathcal{F} by the shape-preserving group $G := \text{Diff}^+(\mathbb{S}^2) \times \text{SO}(3) \rtimes \mathbb{R}^3$: $\mathcal{S} = \mathcal{F}/G$.

It is important to note that the shape space $\mathcal{S} = \mathcal{F}/G$ is a smooth manifold and the canonical projection $\Pi : \mathcal{F} \rightarrow \mathcal{F}/G, f \mapsto [f]$ is a submersion (see for instance [22] and [23]). This submersion is useful in establishing the notion of a vertical space that will be needed a little later. By definition, the vertical space of a submersion is the kernel space of its differential. When the submersion is a quotient map by a group action, the vertical space is the tangent space to the orbit (the terminology comes from the fact that the orbits are usually depicted as vertical fibers over a base manifold which is the quotient space, see Fig.1). In the case of the submersion $\tilde{\Pi} : \mathcal{F} \mapsto \mathcal{F}/\text{Diff}^+(\mathbb{S}^2)$, the vertical space takes a very natural, intuitive form.

Proposition 1: The vertical space $Ver(f)$ of $\tilde{\Pi}$ at some embedding $f \in \mathcal{F}$ is the space of vector fields which are tangent to the shape $f(\mathbb{S}^2)$, or equivalently the space of vector fields such that the dot product

with the unit normal vector field $n_f := \frac{f_u \times f_v}{\|f_u \times f_v\|} : \mathbb{S}^2 \rightarrow \mathbb{R}^3$ vanishes :

$$Ver(f) = \{\delta f : \mathbb{S}^2 \rightarrow \mathbb{R}^3 | \delta f(s) \cdot n_f(s) = 0, \forall s \in \mathbb{S}^2\}.$$

Remark 2: A canonical complement to this vertical space (consisting of tangent vector fields) is given by the space of vector fields normal to the surface $f(\mathbb{S}^2)$ denoted by Nor . This is the sub-bundle of the tangent bundle $T\mathcal{F}$ defined by

$$Nor(f) = \{\delta f : \mathbb{S}^2 \rightarrow \mathbb{R}^3 | \delta f(s) \times n_f(s) = 0, \forall s \in \mathbb{S}^2\}.$$

Any tangent vector $\delta f \in T_f\mathcal{F}$ admits a unique decomposition $\delta f = \delta f^T + \delta f^\perp$ into its tangential part $\delta f^T \in Ver(f)$ and its normal part $\delta f^\perp \in Nor(f)$. Specifically, the normal part is given by:

$$\delta f^\perp = (\delta f \cdot n_f) n_f. \quad (1)$$

See Fig. 2 for an illustration of this decomposition. Generally speaking, one has $T\mathcal{F} = Ver \oplus Nor$ as a direct sum of smooth fiber bundles over \mathcal{F} . This decomposition is preserved by the action of the reparameterization group Γ , i.e. $(\delta f \circ \gamma)^T = \delta f^T \circ \gamma$ and $(\delta f \circ \gamma)^\perp = \delta f^\perp \circ \gamma$ (for a proof of this statement, see Section 1 of the Supplementary Material).

The interest in splitting a perturbation δf into its normal and vertical components comes from the fact that the vertical component $\delta f^T \in Ver(f)$ can only lead to a shape-preserving transformations of the surface $f(\mathbb{S}^2)$. Thus, in the process of deforming one shape into another (for instance along a geodesic path) and quantifying shape differences between them using geodesic lengths, we are not interested in measuring deformations that are in $Ver(f)$. An important novelty of this paper is that the eventual Riemannian metric is imposed only on the δf^\perp components of the perturbations, and that the δf^T components have a zero contribution to the metric.

3 GAUGE INVARIANCE AND RIEMANNIAN METRIC

As mentioned earlier, another important goal of this paper is in developing a framework that is gauge invariant. To appreciate the utility of this framework, we first provide a precise definition and then motivate its use in shape analysis.

3.1 Defining Gauge Invariance

The gauge invariance relates to the parameterization of surfaces along a path in \mathcal{F} and, thus, the mathematical objects of importance in this section are paths $\Psi : [0, 1] \mapsto \mathcal{F}$. The set of such paths is the smooth manifold $\mathcal{P} := \mathcal{C}^\infty([0, 1], \mathcal{F})$.

An element of \mathcal{P} can be thought of as a metamorphosis from the initial shape to the final shape. For instance, Fig. 1 shows two elements in \mathcal{P} as two different deformations from a parameterized horse

to a parameterized cat. To have a picture in mind, consider the upper path $\Psi : t \mapsto \Psi(t)$ in \mathcal{P} : at each time step $t \in [0, 1]$, $\Psi(t)$ is a parameterized shape, i.e. a map from our model manifold \mathbb{S}^2 into \mathbb{R}^3 . The map $\Psi(0)$ is the parameterization of our initial parameterized shape chosen to be a horse and $\Psi(1)$ is the parameterization of our final parameterized shape which, in this case, is a cat.

The definition of length of the path $t \mapsto \Psi(t)$ requires specification of a metric on \mathcal{F} . Given such a metric $(\langle \cdot, \cdot \rangle)$, one can define the length as:

$$L[\Psi] = \int_0^1 (\langle \Psi_t(t), \Psi_t(t) \rangle)_{\Psi(t)}^{\frac{1}{2}} dt, \quad (2)$$

where $\Psi_t(t) = \frac{d\Psi}{dt}(t)$ is the velocity vector of the path $t \mapsto \Psi(t)$, i.e. an infinitesimal deformation of the parameterized shape $\Psi(t)$. The geodesic distance between two shapes f_1 and f_2 is then defined by

$$d(f_1, f_2) = \inf_{\Psi: [0, 1] \rightarrow \mathcal{F} | \Psi(0)=f_1, \Psi(1)=f_2} L[\Psi], \quad (3)$$

where the infimum is taken over all paths connecting shape f_1 and shape f_2 .

We would like the length $L[\Psi]$, for any path Ψ , to match the length of the path $t \mapsto \Psi(t) \circ \gamma(t)$, where $t \mapsto \gamma(t)$ is any time-dependent reparameterization of \mathbb{S}^2 :

$$L[\Psi] = L[\tilde{\Psi}], \quad \text{where } \tilde{\Psi}(t) = \Psi(t) \circ \gamma(t). \quad (4)$$

More formally, set $\Gamma = \text{Diff}^+(\mathbb{S}^2)$ and define the group $\mathcal{G} := \mathcal{C}^\infty([0, 1], \Gamma)$, of time-dependant reparameterizations that acts on \mathcal{P} according to

$$\begin{aligned} \mathcal{G} \times \mathcal{P} &\longrightarrow \mathcal{P} \\ (t \mapsto \gamma(t), t \mapsto \Psi(t)) &\longmapsto (t \mapsto \Psi(t) \circ \gamma(t)). \end{aligned}$$

The group \mathcal{G} is called the *gauge group*, and one says that \mathcal{G} acts by *gauge transformations*. We are looking for a framework where the length of a path is *invariant to gauge transformations*, i.e. satisfies Eqn. (4). One should distinguish these transformations from *temporal* reparameterizations of the path Ψ itself. A gauge transformation changes spatial reparameterization of surfaces, while preserving shapes, along the path, while a temporal reparameterization changes the time it takes to reach each shape along the path.

To build a gauge invariant framework, the basic idea is as follows: take any Γ -invariant Riemannian metric $\langle \langle \cdot, \cdot \rangle \rangle$ on the pre-shape space, and *ignore* the direction tangent to the reparameterization orbit. (An example of Γ -invariant Riemannian metric is the elastic metric defined in Eqn. (7) as it is shown in Section 2 of the Supplementary Material). More precisely, let $\langle \langle \cdot, \cdot \rangle \rangle$ be a Riemannian metric on pre-shape space \mathcal{F} which is preserved by the action of the group of reparameterizations Γ , that is:

$$\langle \langle \delta f_1 \circ \gamma, \delta f_2 \circ \gamma \rangle \rangle_{f \circ \gamma} = \langle \langle \delta f_1, \delta f_2 \rangle \rangle_f, \quad (5)$$

for any $f \in \mathcal{F}$, for any $\delta f_1, \delta f_2 \in T_f \mathcal{F}$ and any $\gamma \in \Gamma$. Given a Γ -invariant sub-bundle H of $T\mathcal{F}$ such that

$$H(f) \oplus Ver(f) = T_f \mathcal{F}, \quad (6)$$

denote by $p_H : T_f \mathcal{F} \rightarrow H(f)$ the projection onto $H(f)$ with respect to the direct sum decomposition given in Eqn. (6). This means that any element $\delta f \in T_f \mathcal{F}$

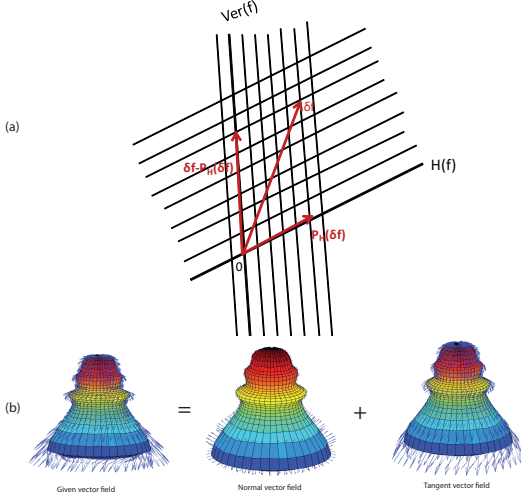


Fig. 2. a. Direct sum decomposition $H(f) \oplus Ver(f) = T_f \mathcal{F}$. b. Vector field decomposition into tangent and normal directions

admits a unique decomposition into the sum of an element $p_H(\delta f) \in H(f)$ and an element in $Ver(f)$. We illustrate this decomposition of vector spaces in Fig. 2.a, while the particular case when H is the space of normal vector fields Nor is shown in Fig. 2.b.

Proposition 3: The non-negative semi-definite inner product on pre-shape space defined by

$$\langle\langle \delta f_1, \delta f_2 \rangle\rangle_f := \langle\langle p_H(\delta f_1), p_H(\delta f_2) \rangle\rangle_f$$

satisfies the gauge invariance condition given in Eqn. (4) and induces a Riemannian metric on quotient space \mathcal{S} such that the quotient map is an isometry between $H(f)$ and the tangent space $T_{[f]} \mathcal{S}$.

3.2 Distinction between Gauge Invariant Framework and Quotient Riemannian Framework

In practice the subbundle H has to be chosen in order to make the implementation easy. A natural choice of subbundle H is the normal bundle Nor which is preserved by the action of the reparameterization group Γ (for a proof of this statement, see Section 1 of the Supplementary Material). We have used this subbundle in the present paper. Another requirement is that the chosen Riemannian metric has to be Γ -invariant. This is the case for the elastic metric defined in next section as we show in Section 2 of the Supplementary Material. We will therefore apply the idea of gauge invariance to the concrete example of the elastic metric and the normal bundle Nor in the remainder of this paper. It is worth noting that

the Riemannian metric on shape space obtained by restricting a Riemannian metric on preshape space to the normal bundle Nor differs in general from the quotient Riemannian metric. In fact, the quotient metric coincides with the restriction to the subbundle Nor if and only if the *Horizontal* subbundle defined by $Hor(f) = \text{Ker}(d\pi)^\perp$ is the normal bundle. This is not the case for the elastic metric. We also remark that the present gauge invariant framework has been used implicitly in [11], Section 6, and [12], Section 11, in the case where the horizontal bundle coincides with the normal bundle.

3.3 Elastic Riemannian Metric

Next, we will choose a Riemannian metric on \mathcal{F} that will enable a gauge-invariant analysis as stated above. We will use the elastic Riemannian metric proposed by Jermyn et al [9] and given in Eqns. (7) and (8). However, before we use this metric we motivate its use by making a connection between the space of parameterized surfaces \mathcal{F} and the space of metrics on a domain, and we will provide some geometrical interpretation of terms in that elastic metric. The space of positive-definite Riemannian metrics on \mathbb{S}^2 will be denoted by $\text{Met}(\mathbb{S}^2)$. Consider a parameterized surface $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$. Denote by $g = f^* \bar{g}$ the pull-back of the Euclidian metric \bar{g} of \mathbb{R}^3 and by n_f the unit normal vector field (Gauss map) on $S = f(\mathbb{S}^2)$.

The metric g and the normal vector field n_f are defined using derivatives of f according to:

$$\begin{aligned} g &= \begin{pmatrix} f_u \cdot f_u & f_u \cdot f_v \\ f_v \cdot f_u & f_v \cdot f_v \end{pmatrix} = \text{Jac}(f)^T \text{Jac}(f), \\ &= \begin{pmatrix} E & F \\ F & G \end{pmatrix}, \quad \text{Jac}(f) = [f_u \ f_v], \quad \text{and} \\ n_f &= \frac{f_u \times f_v}{\|f_u \times f_v\|}, \quad \|f_u \times f_v\| = \sqrt{\det g} = |g|^{\frac{1}{2}}, \end{aligned}$$

where f_u and f_v are the derivatives of f with respect to the local coordinates (u, v) on the sphere. We consider the following relationship between parameterized surfaces on one hand and the product space of metrics and normals on the other :

$$\begin{aligned} \Phi : \mathcal{F} &\longrightarrow \text{Met}(\mathbb{S}^2) \times \mathcal{C}^\infty(\mathbb{S}^2, \mathbb{S}^2) \\ f &\longmapsto (g, n_f). \end{aligned}$$

It follows from the fundamental theorem of surface theory (see Bonnet's Theorem in [24] for the local result, Theorem 3.8.8 in [25] or Theorem 2.8-1 in [26] for the global result) that two parameterized surfaces f_1 and f_2 having the same representation (g, n) differ at most by a translation and rotation. This is an important result, and implies that we can represent a surface by its induced metric $g = f^* \bar{g}$ and the unit normal field $n = n_f$, for the purpose of analyzing its shape. We will not loose any information about the shape of a surface f if we represent it by the pair (g, n) . Let $\delta f_1, \delta f_2$ denote two perturbations of a surface f , and let $(\delta g_1, \delta n_1) = \Phi_*(\delta f_1)$, $(\delta g_2, \delta n_2) = \Phi_*(\delta f_2)$ denote

the corresponding perturbations in (g, n) of f . The expression for Φ_* is given by:

$$\begin{aligned}\delta g &= \text{Jac}(f)^T \text{Jac}(\delta f) + (\text{Jac}(\delta f))^T \text{Jac}(f) \\ &= \begin{pmatrix} 2f_u \cdot \delta f_u & f_u \cdot \delta f_v + f_v \cdot \delta f_u \\ f_u \cdot \delta f_v + f_v \cdot \delta f_u & 2f_v \cdot \delta f_v \end{pmatrix}, \\ \delta n &= -\frac{1}{2} \text{Tr}(g^{-1} \delta g) n + \frac{1}{|g|^{\frac{1}{2}}} (\delta f_u \times f_v + f_u \times \delta f_v).\end{aligned}$$

Then, by definition, the metric on \mathcal{F} used in the present paper measures these perturbations using the expression

$$\begin{aligned}\langle\langle \delta f_1, \delta f_2 \rangle\rangle_f &= \int_{\mathbb{S}^2} ds |g|^{\frac{1}{2}} \{ a \text{Tr}(g^{-1} \delta g_1 g^{-1} \delta g_2) \\ &\quad + \frac{\lambda}{2} \text{Tr}(g^{-1} \delta g_1) \text{Tr}(g^{-1} \delta g_2) + c \delta n_1 \cdot \delta n_2 \}.\end{aligned}\quad (7)$$

The same metric (with $a = 1$) was introduced in [9], Eqn. (2), and called ‘‘elastic metric’’. A related metric measuring the elastic deformation of the interiors of shapes was used in [14] (see Eqn. (4) in [14]). The metric given in Eqn. (7) can be decomposed into three parts

$$\begin{aligned}\langle\langle \delta f_1, \delta f_2 \rangle\rangle_f &= \int_{\mathbb{S}^2} ds |g|^{\frac{1}{2}} \{ a \text{Tr}((g^{-1} \delta g_1)_0 (g^{-1} \delta g_2)_0) \\ &\quad + b \text{Tr}(g^{-1} \delta g_1) \text{Tr}(g^{-1} \delta g_2) + c \delta n_1 \cdot \delta n_2 \},\end{aligned}\quad (8)$$

where $b = \frac{\lambda+a}{2}$ and where A_0 is the traceless part of a 2×2 -matrix A defined as $A_0 = A - \frac{\text{Tr}(A)}{2} I_{2 \times 2}$. The term multiplied by a measures area-preserving changes in the induced metric g , the term multiplied by b measures changes in the area of patches, and the last term measures bending. Note that only the relative weights b/a and c/a are meaningful.

Now we consider a key property of this metric that relates to reparameterization of surfaces. Recall that $\Gamma := \text{Diff}^+(\mathbb{S}^2)$ denotes the subgroup of $\text{Diff}(\mathbb{S}^2)$ consisting of diffeomorphisms γ which preserve the orientation of \mathbb{S}^2 , i.e. such that $\det \text{Jac}(\gamma) > 0$. (Note that for a diffeomorphism $\gamma \in \text{Diff}(\mathbb{S}^2)$, since $\text{Jac}(\gamma)$ is invertible, the determinant of $\text{Jac}(\gamma)$ never vanish. It follows that either $\det \text{Jac}(\gamma)(s) > 0$ for all $s \in \mathbb{S}^2$, or $\det \text{Jac}(\gamma)(s) < 0$ for all $s \in \mathbb{S}^2$.) It will be called the group of orientation-preserving reparameterizations. The group $\Gamma = \text{Diff}^+(\mathbb{S}^2)$ acts on $\text{Maps}(\mathbb{S}^2, \mathbb{R}^3)$ by pre-composition. That is, a surface f is reparameterized by a $\gamma \in \text{Diff}^+(\mathbb{S}^2)$ according to $f \mapsto f \circ \gamma^{-1}$. How does the metric-normal representation (g, n) of that surface change due to reparameterization? This representation of the reparameterized surface is given by $(\gamma^{-1*} g, n \circ \gamma^{-1})$. This representation is Γ -equivariant for the actions introduced, i.e. if we reparameterize a surface and then compute its (g, n) representation, or if we compute (g, n) representation of a surface and then reparameterize them according to $(\gamma^* g, n \circ \gamma)$, we get the same result.

Proposition 4: The elastic metric is invariant to the action of $\text{Diff}^+(\mathbb{S}^2)$.

Proof: Please refer to Section 2 of the Supplementary Material.

Although this elastic metric has been introduced by Jermyn et al. [9], it has not been used completely for shape analysis of surfaces. Furthermore, we are going to use it in a novel way – by restricting its evaluation only to the normal vector fields on a surface (see next section for a geometric expression of the resulting metric on shape space).

Definition 2: For any two perturbations $\delta f_1, \delta f_2 \in T_f \mathcal{F}$ define the pairing

$$\langle\langle \delta f_1, \delta f_2 \rangle\rangle_f = \langle\langle \delta f_1^\perp, \delta f_2^\perp \rangle\rangle_f,$$

where δf_i^\perp is the normal component of δf_i as defined in Eqn. (1) and where $\langle\langle \cdot, \cdot \rangle\rangle_f$ is as given in Eqn. (8).

Remark 5: It follows from proposition 3, that $(\langle\langle \cdot, \cdot \rangle\rangle)$ satisfies the gauge-invariant condition $L[\Psi] = L[\tilde{\Psi}]$, where Ψ is any path of shapes, $\tilde{\Psi}(t) = \Psi(t) \circ \gamma(t)$ with $t \mapsto \gamma(t)$ any time-dependant reparameterization, and $L[\Psi]$ is as specified in Eqn. (2).

3.4 Geometric expression of the elastic metric in the normal direction

In this section, we will give some geometric interpretation of the restriction of the elastic metric on the space of normal vector fields introduced in the previous section. Given a surface f parameterized by (u, v) , we will consider normal variations:

$$f_\varepsilon(u, v) = f(u, v) + \varepsilon h(u, v) n(u, v),$$

where $(u, v) \in \mathbb{S}^2$, $\varepsilon > 0$, $n(u, v) = n_f(u, v)$ is the unit normal to the surface $f(\mathbb{S}^2)$ at $f(u, v)$, and $h : \mathbb{S}^2 \rightarrow \mathbb{R}$ is a real function corresponding to the amplitude of the normal vector field $h n_f$. Let us compute the first fundamental form g_ε of the surface parameterized by f_ε , i.e. the metric induced on the parameterized surface f_ε by the Euclidian metric of \mathbb{R}^3 . We obtain

$$f_{\varepsilon,u} := \frac{\partial f_\varepsilon}{\partial u} = f_u + \varepsilon h n_u + \varepsilon h_u n, \quad (9)$$

$$f_{\varepsilon,v} := \frac{\partial f_\varepsilon}{\partial v} = f_v + \varepsilon h n_v + \varepsilon h_v n.$$

Therefore

$$f_{\varepsilon,u} \cdot f_{\varepsilon,u} = f_u \cdot f_u + 2\varepsilon h n_u \cdot f_u + \varepsilon^2 (h^2 n_u \cdot n_u + h_u^2),$$

where we have used that $n \cdot f_u = 0$ and $n_u \cdot n = 0$ since $n \cdot n = 1$. Similarly

$$f_{\varepsilon,v} \cdot f_{\varepsilon,v} = f_v \cdot f_v + 2\varepsilon h n_v \cdot f_v + \varepsilon^2 (h^2 n_v \cdot n_v + h_v^2),$$

and

$$f_{\varepsilon,u} \cdot f_{\varepsilon,v} = f_u \cdot f_v + \varepsilon h (n_u \cdot f_v + f_u \cdot n_v) + \varepsilon^2 (h^2 n_u \cdot n_v + h_u h_v).$$

It follows that

$$\begin{aligned}g_\varepsilon &= g + 2\varepsilon h \begin{pmatrix} n_u \cdot f_u & n_u \cdot f_v \\ n_v \cdot f_u & n_v \cdot f_v \end{pmatrix} \\ &\quad + \varepsilon^2 h^2 \begin{pmatrix} n_u \cdot n_u & n_u \cdot n_v \\ n_v \cdot n_u & n_v \cdot n_v \end{pmatrix} + \varepsilon^2 \begin{pmatrix} h_u^2 & h_u h_v \\ h_u h_v & h_v^2 \end{pmatrix}.\end{aligned}$$

Using the definition of the second fundamental form $\mathbf{\Pi}$ of the surface $f(\mathbb{S}^2)$, we obtain

$$g_\varepsilon = g - 2\varepsilon h \mathbf{\Pi} + \varepsilon^2 h^2 \mathbf{\Pi} g^{-1} \mathbf{\Pi} + \varepsilon^2 \begin{pmatrix} h_u & h_v \end{pmatrix}^T \begin{pmatrix} h_u & h_v \end{pmatrix}.$$

It follows that

$$g^{-1} \delta g = -2h g^{-1} \mathbf{\Pi} = -2h \mathbf{L}, \quad (10)$$

where \mathbf{L} is called the *shape operator*. Recall that the eigenvalues of \mathbf{L} are the principal curvatures of the surface $f(\mathbb{S}^2)$, denoted by κ_1 and κ_2 , which provide local information about the surface: at a given point on the surface, they measure the greatest and smallest possible curvatures of a curve drawn on the surface passing through this point. For instance, the vanishing of the principal curvatures at one point of the surface tells that the surface is flat near this point (i.e. looks like a plane). The equality $\kappa_1 = \kappa_2 = 1/R$ at one point tells that the surface looks like a sphere of radius R near this point. In other words, κ_1 and κ_2 are functions on the surface that characterize how the surface is locally curved.

On the other hand, the variation δn of the normal vector field satisfies $\delta n \cdot n = 0$ since the norm of n remains constant. Moreover $n \cdot f_u = n \cdot f_v = 0$, therefore $\delta n \cdot f_u = -n \cdot \delta f_u$ and $\delta n \cdot f_v = -n \cdot \delta f_v$. By Eqn. (9), $\delta f_u = h n_u + h_u n$, hence $\delta n \cdot f_u = -h_u$ and similarly $\delta n \cdot f_v = -h_v$. Consequently $\delta n = \alpha f_u + \beta f_v$ where $\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = -g^{-1} \begin{pmatrix} h_u \\ h_v \end{pmatrix}$. It follows that for two normal vector fields hn and kn with $h, k \in C^\infty(\mathbb{S}^2, \mathbb{R})$, one has

$$\delta n_1 \cdot \delta n_2 = \begin{pmatrix} h_u & h_v \end{pmatrix} g^{-1} \begin{pmatrix} k_u \\ k_v \end{pmatrix}. \quad (11)$$

Using Eqn. (10) and Eqn. (11) the elastic metric restricted to these normal fields is given by :

$$\begin{aligned} ((hn, kn))_f &= \int_{\mathbb{S}^2} ds |g|^{\frac{1}{2}} \{ h k (2a(\kappa_1 - \kappa_2)^2 \\ &+ 4b(\kappa_1 + \kappa_2)^2) + c \begin{pmatrix} h_u & h_v \end{pmatrix} g^{-1} \begin{pmatrix} k_u \\ k_v \end{pmatrix} \} \end{aligned} \quad (12)$$

This is the form used to define and compute geodesic paths in the shape space \mathcal{S} in this paper. The difference $\kappa_1 - \kappa_2$ in the first term has been called the normal deformation of the surface in [27]. The sum $\kappa_1 + \kappa_2$ is twice the mean curvature which measures variations of the area of local patches. These two terms are related to the shape index $\text{idx} = \frac{2}{\pi} \arctan \frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}$ [28]. The last term in Eqn. (12) measures variations of the normal vector field, i.e. bending.

4 GEODESIC COMPUTATION

Finding geodesics between two surfaces f_1 and f_2 under invariant Riemannian metrics is a difficult problem. In the present case, analytical solutions are not known and we will use a path-straightening approach to find geodesics. This method has been used for instance in [9] and [7]. The basic idea here is to connect f_1 and f_2 by any initial path and then iteratively straighten it until it becomes a geodesic. The update is performed using the gradient of an energy function.

As mentioned earlier, this method only achieves a local minimum of the energy function, resulting in a geodesic path that may not be the shortest geodesic.

4.1 Removing rotations and translations

Since we are only interested in shapes of objects and not in the way objects are oriented or placed in the ambient space \mathbb{R}^3 , we have to remove the actions of the rotation and translation groups. In theory, we could deal with them as we do with the group of reparameterizations. However, since $\text{SO}(3) \times \mathbb{R}^3$ is just a 6-dimensional Lie group in comparison to the infinite-dimensional Fréchet Lie group $\text{Diff}^+(\mathbb{S}^2)$, it is more efficient to do the following. First find the best translation and rotation that align the two objects to be compared, and then find the geodesic between them. To center an object we use Algorithm 3 given in Section 3 of the Supplementary Material to compute the center of mass and then subtract it from the surface coordinates.

There are many ways to find the best rotation that align two shapes. In the case of elongated objects (which was the case in our applications), one can do the following. Given two shapes S_1 and S_2 , find the best ellipsoids \mathbb{E}_1 and \mathbb{E}_2 that approximate the cloud of points defining S_1 and S_2 respectively, and the unitary matrices U_1 and U_2 that map the reference axes to the axes of the ellipsoids (with decreasing lengths). The unitary matrices U_1 and U_2 are uniquely defined if the approximating ellipsoids are triaxial (i.e. the lengths of their principal axes are distinct). Then, we can apply the product matrix $U_2 U_1^{-1}$ on the shape S_1 to rotationally align with S_2 . If one encounters a 180 degree flip, apply instead $U_2 R U_2^{-1}$, where R is the 180 degree rotation around the z -axis. As an example, Fig. 3 shows two hands that have different orientations in space, the corresponding ellipsoids, and the hands after rotation (with a gap to separate them in order to facilitate visualisation).




Fig. 3. Rotational alignment: two hands before and after the alignment, respectively at the left and at the right. Each hand is approximated by an ellipsoid. The rotation used apply the axis of one ellipsoid to the axis of the other.

To find the best ellipsoid that approximate a surface S and the corresponding rotation U , one can use a singular value decomposition of $S^T S$. However, in the case where the surface is the boundary of a 3D-volume, it is more accurate to compute the mean of $S^T S$ over the inscribed volume. It has also a more physical meaning since the resulting ellipsoid is equivariant with respect to affine transformations (see Section 3 of the Supplementary Material where

the effect of the rotations of the initial surface on the ellipsoid is illustrated (Fig. 2) and where detailed Algorithms are presented). Moreover, the estimation of ellipsoid for an inscribed volume is more stable under reparameterizations. To illustrate this robustness we show in Fig. 4 different parameterizations of a horse (middle row) obtained by pre-composing a given parameterization by a diffeomorphism of the sphere (bottom row) and the resulting ellipsoid. The diffeomorphisms used in this experiment are (from left to right) $\varphi_1 = \text{identity}$, $\varphi_2 = \text{rotation of } -3\pi/4 \text{ around } x\text{-axis}$, $\varphi_3 = \text{Möbius transformation that maps } z \in \mathbb{S}^2 \simeq \mathbb{C} \cup \{\infty\} \text{ to } \phi_3(z) = 0.4z + 0.5$, $\varphi_4 = \text{rotation of } -\pi/2 \text{ around } x\text{-axis composed with } \varphi_3$.

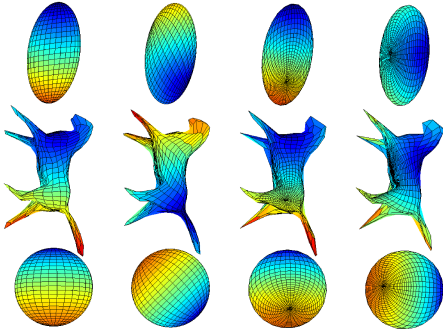


Fig. 4. Robustness of the approximating ellipsoid of a surface with respect to reparameterizations.

In the case where the approximating ellipsoids are not triaxial, one has to use additional information about the surfaces to align them properly (for instance, one can use four points on each surface). This case was not implemented in the present paper.

4.2 Computations of the energy

Let $\Psi : [0, 1] \rightarrow \mathcal{F}$. The energy of the path Ψ is defined to be:

$$\mathcal{E}(\Psi(t)) = \int_0^1 \langle \Psi_t^\perp, \Psi_t^\perp \rangle_{\Psi(t)} dt = \int_0^1 \langle (\Psi_t, \Psi_t) \rangle_{\Psi(t)} dt,$$

where $\langle \cdot, \cdot \rangle$ is the elastic metric given in Eqn. (8), $\Psi_t^\perp = (\Psi_t \cdot n)n$ is the normal component of the deformation, and $\langle (\cdot, \cdot) \rangle$ is the inner product presented in Eqn. (12). We will present several numerical strategies for approximating this energy and will compare their computational costs in Table 1. This evaluation uses a linear path connecting two concentric spheres of radius $R_1 = 1$ and $R_2 = 2.5$, with constants $a = 1$, $\lambda = 0.125$ and $c = 0$ for defining energy (see Fig. 5). The theoretical value of the energy in this case is given by $E_{th} = 32\pi(a + \lambda)(R_2 - R_1)^2$ and measures exclusively the cost of changing the area of the spheres (the first and third term of the metric given in Eqn. (8) vanish in this experiment). We expect that improvement in accuracy comes at an increased computational cost, and this is indeed the

case in the results presented in the Table. Note that a time-dependent rotation is applied on the path of spheres, but the values of the energy is independant of this rotation.




Fig. 5. Path connecting two concentric spheres used for computations in Table 1.

One way to compute the energy of a path Ψ of shapes is to express it using the coefficients of the first fundamental form. Consider the mapping $\Psi : \mathbb{S}^2 \times \mathbb{R} \rightarrow \mathbb{R}^3$ and define

$$E = \Psi_u \cdot \Psi_u, F = \Psi_u \cdot \Psi_v, G = \Psi_v \cdot \Psi_v, \quad (13)$$

and their time derivatives

$$\dot{E} = 2\Psi_{tu}^\perp \cdot \Psi_u, \dot{F} = \Psi_{tu}^\perp \cdot \Psi_v + \Psi_u \cdot \Psi_{tv}^\perp, \dot{G} = 2\Psi_{tv}^\perp \cdot \Psi_v,$$

as well as the unit normal field $n := n_f = \frac{f_u \times f_v}{\|f_u \times f_v\|}$ and the vector field $w = \Psi_{tu}^\perp \times \Psi_v + \Psi_u \times \Psi_{tv}^\perp$. Then, the energy of a path Ψ decomposes into the sum of four terms: $\mathcal{E}(\Psi(t)) = \mathcal{E}_1 + \mathcal{E}_2 + \mathcal{E}_3 + \mathcal{E}_4$, where

$$\begin{aligned} \mathcal{E}_1 &= a \int_0^1 \int_{\mathbb{S}^2} (EG - F^2)^{-3/2} B \, du \, dv \, dt \\ &\quad \text{with } B = G^2 \dot{E}^2 + 2(EG + F^2) \dot{F}^2 + E^2 \dot{G}^2 \\ &\quad \quad - 4FG \dot{E} \dot{F} + 2F^2 \dot{E} \dot{G} - 4EF \dot{F} \dot{G}, \\ \mathcal{E}_2 &= \left(\frac{\lambda}{2} + \frac{c}{4}\right) \int_0^1 \int_{\mathbb{S}^2} (EG - F^2)^{-3/2} (G \dot{E} - 2F \dot{F} + E \dot{G})^2 \, du \, dv \, dt, \\ \mathcal{E}_3 &= -c \int_0^1 \int_{\mathbb{S}^2} (G \dot{E} - 2F \dot{F} + E \dot{G}) (n \cdot w) \, du \, dv \, dt, \\ \mathcal{E}_4 &= c \int_0^1 \int_{\mathbb{S}^2} (EG - F^2)^{-1/2} (w \cdot w) \, du \, dv \, dt. \end{aligned}$$

In the implementation of these formulas, we can reach singularities on the boundary of the integration domain, which we can ignore. In the example involving two concentric spheres, the total energy computed by this method is labelled $E_{I\&II}$ in Table 1.

Another way to compute the energy is based on Eqn. (12) that expresses the elastic metric in terms of principal curvatures. In terms of the coefficients of the first fundamental form given in Eqn. (13) and of the second fundamental given by

$$\begin{aligned} e &= \Psi_{uu} \cdot n = -\Psi_u \cdot n_u, \\ f &= \Psi_{uv} \cdot n = -\Psi_u \cdot n_v = -\Psi_v \cdot n_u, \\ g &= \Psi_{vv} \cdot n = -\Psi_v \cdot n_v, \end{aligned}$$

the Gauss curvature K and the mean curvature H have the following expressions

$$K = \frac{eg - f^2}{EG - F^2}, \quad H = \frac{1}{2} \frac{eG + gE - 2fF}{EG - F^2},$$

and the principal curvatures are given by

$$\kappa_1 = H + \sqrt{H^2 - K}, \quad \kappa_2 = H - \sqrt{H^2 - K}.$$

Again, in the implementation of these formulas, we can get singularities for curvatures on the boundary, but we can ignore them in computing the integral

given in Eqn. (12). This corresponds to cutting a small disc on the parameterized surface around the images of the north and south poles.

In the example of the two concentric spheres, the theoretical values of κ_1 and κ_2 is the constant function equal to $1/R$ where $R = R_1 + t(R_2 - R_1)$ is the radius of the sphere along the path interpolating linearly the sphere of radius $R_1 = 1$ to the sphere of radius $R_2 = 2.5$. The total energy computed by this method is labelled $E_{k_1 k_2}$ in Table 1.




Fig. 6. From left to right: A hand with the tangent plane and normal at the tip of the index finger; 3-neighborhood of the tip of the index finger; tip of the index finger after rotation; a closeup; approximating second order polynomial.

To improve the computation of the curvatures and therefore also of the energy, we can use polynomial approximations of the surfaces. This procedure, leading to the computation of the principal curvatures, is illustrated in Fig. 6. To compute the principal curvatures at a given point of a surface, e.g. at the tip of the index finger of the hand depicted in Fig. 6, we first compute the normal at this point by averaging the normals of the facets having this point as vertex. A tangent plane is then defined as the plane orthogonal to the normal passing through the point under consideration. A neighborhood of the point is isolated from the surface (we use a 3-neighborhood, see second drawing in Fig. 6). We then apply a rigid transformation to center the point at the origin and to align the tangent plan with the xy -plane (see third drawing, and a closeup in the fourth drawing). After that, we use Algorithm 5 given in Section 5 of the Supplementary Material to compute the second order polynomial $P(x, y) = a_1x^2 + a_2y^2 + a_3xy + a_4x + a_5y + a_6$, which minimizes the sum $\sum_i (z_i - P(x_i, y_i))^2$ over the points of the centered and rotated neighborhood. Then, the Gauss curvature at that point is given by $K = 4a_1a_2 - a_3^2$, the mean curvature by $H = a_1 + a_2$, and the principal curvatures by $\kappa_1 = a_1 + a_2 + \sqrt{((a_1 - a_2)^2 + a_3^2)}$ and $\kappa_2 = a_1 + a_2 - \sqrt{((a_1 - a_2)^2 + a_3^2)}$. In the example of the two concentric spheres, the total energy computed using the principal curvatures obtained by this method is labelled E_P in Table 1.

In order to show that the energy function of a path of shapes is independent of the way the objects are parameterized, we replace the integration over the domain of parameterization by the integration over the triangulated surfaces. This means that we approximate the area elements of the surfaces by the area of triangles whose vertices are given by the parameterization. In this way, the parameterisation




Fig. 7. A path of zero energy connecting a hand and the same hand with another parameterization.

Energy, 10^4 points per object	Elapsed time for 10^4 points
$E_{I \& II} = 246.2854$	0.221726 seconds
$E_{k_1 k_2} = 249.1969$	0.862376 seconds
$E_P = 255.8288$	1.238354 seconds
$E_\Delta = 255.9043$	9.738431 seconds
Energy for 4×10^4 points	Elapsed time for 4×10^4 points
$E_{I \& II} = 249.1503$	0.978828 seconds
$E_{k_1 k_2} = 251.8494$	3.45599 seconds
$E_P = 254.7646$	4.906798 seconds
$E_\Delta = 254.7832$	39.011899 seconds

TABLE 1

Computation of the energy of a path connecting two concentric spheres (Fig. 5) using different methods, and time needed for the computations. The theoretical value of the energy is $E_{th} = 254.4690$. Here $R_1 = 1$, $R_2 = 2.5$, $\lambda = 0.125$ and $c = 0$.

of surfaces is only used to define the surfaces, but play no role at all in the computation of the energy function. In the example of the two concentric spheres, the total energy computed by this method is labelled E_Δ and given in Table 1. In Fig. 7, a path connecting a hand to the same hand but with a different parameterization is shown. The energy of this path, computed with the constants $a = 1$, $\lambda = c = 0.125$, reads $E_\Delta = 0.4824$, hence is close to 0. Now returning to Fig. 1, the energy of the lower path from a horse to a cat computed with the same constants is $E_\Delta = 227.4049$, its length is $L[\Psi] = 14.9099$, whereas the upper path (obtained from the lower path by applying a different reparameterization at each time step) has an energy equal to $E_\Delta = 225.5249$ and a length of $L[\Psi] = 14.8802$. Note that in this example, the colors refer to the Euclidean distance to the point on the surface corresponding to the image of the north pole of the sphere (cold colors for small distances versus hot colors for large distances). In particular, the north and south poles do not correspond in these two paths.

4.3 Orthonormal Basis of Deformations

In this section, we define bases for representing perturbations of a path of surfaces. These basis elements form possible directions for use in path-straightening in Section 4.4. The first basis we used is a variation

of the one given in [7]. We start with a basis $\mathcal{B}_1 = \{Y_l^m, 1 \leq l \leq N, -l \leq m \leq l\}$ of spherical harmonics of degree less than N , available in Matlab as function SPHARM (see [29] for more information on spherical harmonics). We make three copies of this basis of \mathbb{R} -valued functions in order to obtain a basis \mathcal{B}_2 of the space $L^2(\mathbb{S}^2, \mathbb{R}^3)$ of \mathbb{R}^3 -valued functions. Similar to Xie et al. [30], we demonstrate reconstruction of some surfaces using the resulting basis, as the degree of the spherical harmonics grows, in the Supplementary Material (Fig. 5).

Next, we want to construct a basis of perturbations of a path connecting two parameterized surfaces f_1 and f_2 . In order to apply the path-straightening method as described in Section 4.4, we want the perturbations to vanish at $t = 0$ and $t = 1$ so that f_1 and f_2 remain fixed. Therefore, we want a basis of $L^2(\mathbb{S}^2 \times [0, 1], \mathbb{R}^3)$ with elements that have this property. To ensure this, each element of \mathcal{B}_2 is multiplied by a basis element of $L^2([0, 1], \mathbb{R})$ of the form $P_j(t) = \frac{1}{4} \sin(\pi j t)$, $1 \leq j \leq J$. Unfortunately a major limitation of the resulting L^2 basis is that slowly- and rapidly-oscillating harmonics have comparable amplitudes. In the implementation of the path-straightening method, this implies that the updated path can go out of the open set of immersions.

One possible way to counter this effect is to orthonormalize the L^2 -basis with respect to an H^1 -type scalar product (i.e. that measures also the variation of the derivatives). For this kind of scalar product, an orthonormal basis consists of functions which have controlled derivatives (hence can not oscillate too much). This approach was also used in [7] where the L^2 -basis is orthonormalized using the Gram-Schmidt procedure with respect to the following scalar product

$$(B^1, B^2) = \int_0^1 \int_{\mathbb{S}^2} (B^1 \cdot B^2 + B_t^1 \cdot B_t^2 + B_u^1 \cdot B_u^2 + B_v^1 \cdot B_v^2 + B_{t,u}^1 \cdot B_{t,u}^2 + B_{t,v}^1 \cdot B_{t,v}^2) ds dt.$$

However, when increasing the degree of spherical harmonics, the computational cost of generation of an orthonormal basis using this scalar product is very high. Therefore, we first orthonormalize the basis \mathcal{B}_2 with respect to the following inner product

$$(B^1, B^2) = \int_{\mathbb{S}^2} (B^1 \cdot B^2 + B_u^1 \cdot B_u^2 + B_v^1 \cdot B_v^2) ds, \quad (14)$$

and then we multiply the resulting basis by the time-dependant components $P_j(t) = \frac{1}{4} \sin(\pi j t)$, $1 \leq j \leq J$. The advantage of this method is that the Gram-Schmidt procedure is applied to matrices of lower dimensions (without the time dimension) and on a smaller number of elements (by a factor J). The spatial oscillations of the resulting basis elements are well controlled by the presence of the spatial derivatives B_u and B_v in the inner product given in Eqn. (14).

4.4 Path-straightening method

The path-straightening method is used to find critical points of the energy functional. Starting with an

arbitrary path, the method consists of iteratively deforming (or ‘‘straightening’’) the path in the opposite direction of the gradient, until the path converges to a geodesic. The gradient of the path energy is approximated using a basis \mathcal{B} of possible perturbations of a path of surfaces Ψ , as constructed in the previous section. We first compute the directional derivatives $\nabla \mathcal{E}_\Psi(b) = \frac{d}{de}(\mathcal{E}(\Psi + \epsilon b))|_{\epsilon=0}$ where b ranges over \mathcal{B} . This is done by fixing a small ϵ_1 and approximating the directional derivative by $\nabla \mathcal{E}_\Psi(b) \simeq (\mathcal{E}(\Psi + \epsilon_1 b) - \mathcal{E}(\Psi))\epsilon_1^{-1}$. Using the finite orthonormal basis \mathcal{B} , we obtain a numerical approximation of the gradient: $\nabla \mathcal{E}_\Psi = \sum_{b \in \mathcal{B}} \nabla \mathcal{E}_\Psi(b) \cdot b$. In particular, the norm of the gradient is approximately given by $\|\nabla \mathcal{E}_\Psi\|^2 = \sum_{b \in \mathcal{B}} \nabla \mathcal{E}_\Psi(b)^2$. The update of the path is done by replacing Ψ by $\Psi - \epsilon_2 \nabla \mathcal{E}_\Psi$, where ϵ_2 is a small parameter that has to be adjusted empirically. The method is detailed in Algorithm 1 below.

Input:

- 1) A path Ψ between two parameterized surfaces f_1 and f_2 ,
- 2) a basis of perturbation \mathcal{B} .

Output:

- 1) The minimal energy needed to deform f_1 into f_2 given by the value of the cost function E ,
- 2) the geodesic path between f_1 and f_2 .

Set $\|\nabla E\|^2 = 1$.

while $\|\nabla E\|^2 > 10^{-3}$ **do**

2- Compute the energy E of the path Ψ according to Eqn.(8) or Eqn. (12).

3- Set $\Psi_{\text{upd}} = 0$ and $\|\nabla E\|^2 = 0$.

for $i \leftarrow 1$ **to** $\text{size}(\mathcal{B})$ **do**

4- Add a perturbation to the current path Ψ : define $\Psi(i) = \Psi + \epsilon_1 \mathcal{B}(i)$, where $\mathcal{B}(i)$ is the element of the perturbation basis \mathcal{B} of index i and $\epsilon_1 > 0$ is small.

5- Compute the energy $E(i)$ of the perturbed path $\Psi(i)$.

6- Compute the gradient of energy $\nabla E(i)$ in the direction $\mathcal{B}(i)$ using the approximation $\nabla E(i) \sim \frac{E(i) - E}{\epsilon_1}$.

7- Compute the updating path:

$\Psi_{\text{upd}} \leftarrow \Psi_{\text{upd}} + \nabla E(i) \cdot \mathcal{B}(i)$.

8- Compute the squared norm of the gradient of energy at path Ψ : $\|\nabla E\|^2 \leftarrow \|\nabla E\|^2 + (\nabla E(i))^2$.

end

10- Update the path: $\Psi = \Psi - \epsilon_2 \Psi_{\text{upd}}$

end

Algorithm 1: Path-straightening method.

5 EXPERIMENTAL RESULTS

The 3D realistic models used in our experiments are part of the TOSCA [31] dataset. Their spherical parameterizations were initially implemented in [32].

5.1 Examples of geodesics obtained by path-straightening

First we apply the path-straightening method to the case where the surfaces at the extremes of the initial path have the same shape, but different parameterizations. More precisely, we consider the special case where $\Psi_0(0) = f_1$, $\Psi_0(1) = f_1 \circ \gamma$ for some diffeomorphism γ and where we initialize the path with piecewise linear interpolation to a different surface f_3 in the middle of the path, i.e. $\Psi_0(\frac{1}{2}) = f_3$. This situation is illustrated in Fig. 8. The proposed gauge-invariant approach is expected to reach a path with constant shape as a geodesic, despite the different

shapes appearing in the initial path and the different parameterization of shapes at the end points of the path (to emphasize the differences in parameterization, zoom-ins of these surfaces are also shown). Once we have the geodesic path Ψ between the given surfaces, the distance in the shape space between f_1 and $f_1 \circ \gamma$, $d_\Psi(f_1, f_2)$, is the length of Ψ as specified in Eqn. (2). As expected, the resulting geodesic path, shown in Fig. 8, is constant with the same shape as the either end, and with $d_\Psi(f_1, f_2) = 0$. Using path-straightening, we obtain a 99.28% decrease in the energy function from the initial path to the final path.

In Fig. 9 we consider more challenging shapes. The top two rows corresponds to the case where we have $\Psi_0(0) = f_1$, $\Psi_0(1) = f_1$ (a cat) and where we initialize the path with piecewise linear interpolation to a horse in the middle of the path. The upper row shows the initial path and the second row the geodesic path. We can see that the geodesic path has a constant shape throughout, as expected. We also plot the evolution of the path energy on the right during path-straightening. We can see that the energy decreases until it reaches a relatively small value; the theoretical minimum is, of course, zero for a constant path. In the last two rows of Fig. 9, we consider the case of two hands. We initialize the path with linear interpolation (third row in Fig. 9), and the resulting path is shown in the last rows of Fig. 9. The energy evolution is shown on the right and we can see the energy decreasing until it reaches a constant value; thus, the final path is a geodesic. It can be seen that the deformation along the geodesic path is more natural than the original path.




Fig. 8. Illustration of initial path (upper row) and geodesic path in shape space (middle row). The energy is reported in the bottom row. The surfaces at the end points of the path have different parameterizations.

5.2 Classification of 3D shapes

As mentioned earlier, the geodesic paths provide us with tools for comparing, and deforming parameterized surfaces. We suggest a comparison of shapes of 3D objects using geodesic distances between their boundary surfaces in the shape space. This section presents a specific application to illustrate that idea. In this section, we study several shapes belonging to four classes: horses, hands, cats and centaurs.

We begin by computing the pairwise geodesic distances between corresponding 3D surfaces. The distance matrix and the classification dendrogram are shown in Fig. 10. In the distance matrix, we can easily distinguish four classes corresponding to four blue boxes. Actually the cold colors in the illustrated matrix correspond to small values of distances versus hot colors that correspond to greater distances. The clustering obtained using the *dendrogram* (command in matlab) can be interpreted by slicing the top of the dendrogram by a horizontal line to split the shapes into the desired number of classes, and then sliding the horizontal line to the bottom in order to refine the classification. The coarsest classification results by slicing the dendrogram into two classes (by a horizontal line close to the top), the shapes 4, 5 and 6 (the hands) forms a first class and the remaining (horses, cats and centaurs) are grouped together as a second class. The next level in classification distinguishes the shapes 1, 2, and 3 (the horses) and 12, 13 (the centaurs) from the shapes 7, 8, 9, 10, 11 (the cats). The finest level separates the horses and the centaurs in different classes and results in four classes. Thus, we argue that the proposed framework provides a powerful tool for shape classification.

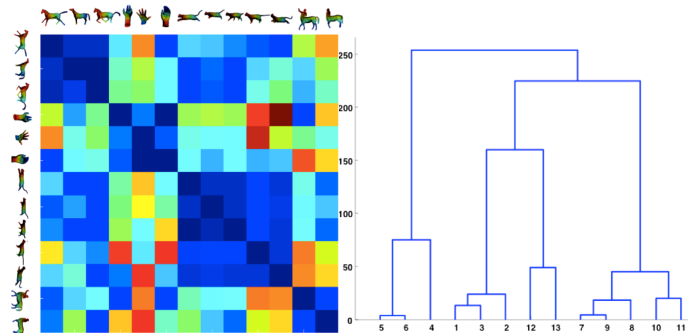


Fig. 10. Classification performance; left: the distance matrix. right: the dendrogram.

5.3 The effect of number of basis elements

In this section, we study the effect of the number of basis elements, used in path-straightening, on the resulting geodesic path. Given two parameterized surfaces f_1 and f_2 , we again initialize the path with the linear interpolation to a different surface f_3 in the middle of the path. This initial path is shown in the upper row of Fig. 11. Then, we compute the geodesic path using different number of basis elements. We show the geodesic paths that use, respectively, 52, 432 and 1728 basis elements. We can see that the larger the number of basis elements, the better the final result is. We also provide the trade-off between the number of basis elements and the minimum energy value obtained. The trade-off confirms our assertion. At the bottom of the figure, we show the geodesic path obtained when the path-straightening Algorithm

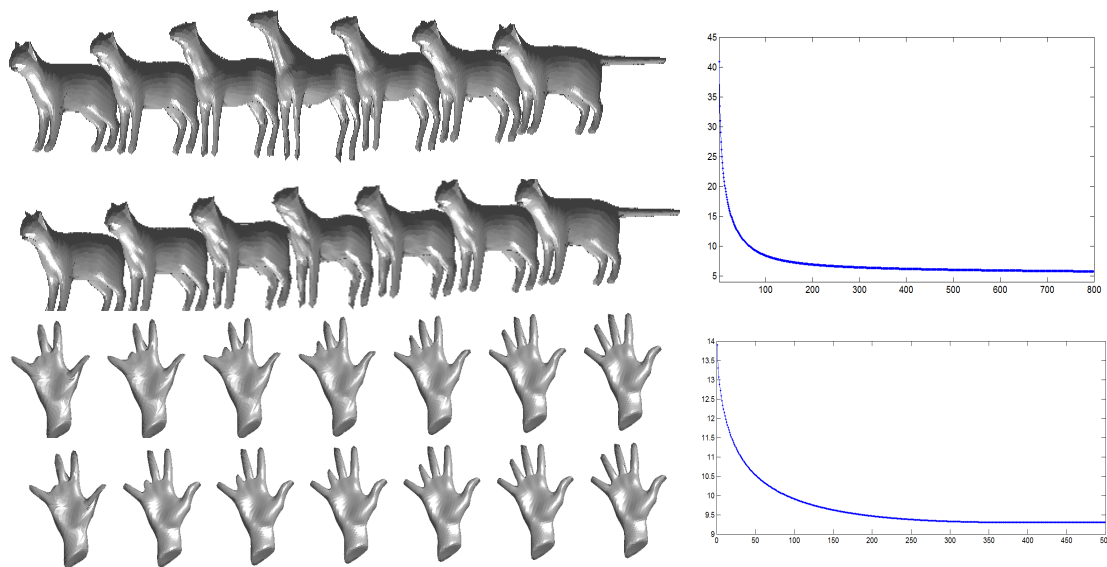


Fig. 9. The top row shows an initial path formed by linear interpolation between a cat to a horse and back to the cat. The second row illustrates the geodesic obtained after 800 iterations of path-straightening. The corresponding evolution of the energy is shown on the right. Similarly, the third row shows a linear path between two hands with bad correspondence and the last row shows the final geodesic, with the corresponding energy is shown on the right.

is initialized with the linear interpolation between f_1 and f_2 . This path is also calculated using the number of basis elements corresponding to the lowest energy. This path can be seen as ground truth to visually interpret the previous geodesics (with more complicated initial conditions and fewer basis elements).

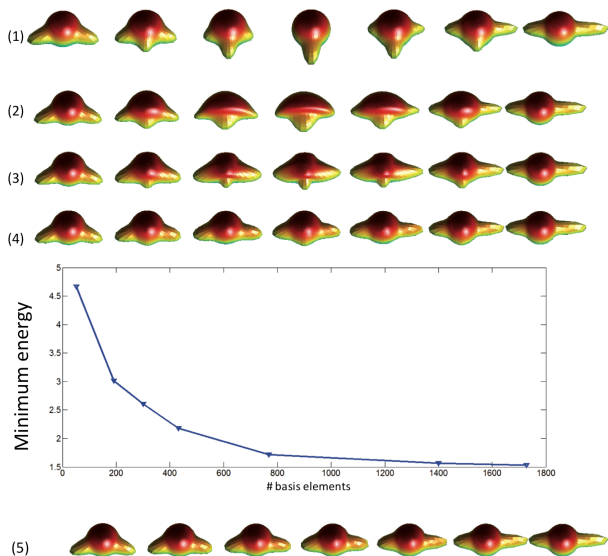


Fig. 11. The effect of the number of basis elements, (1) initial path, (2) geodesic path using 52 basis elements, (3) geodesic path using 432 basis elements, (4) geodesic path using 1728 basis elements, (5) geodesic path using 1728 basis elements after linear interpolation initialization.

6 CONCLUSION

In this paper we have proposed a novel Riemannian framework for computing geodesic paths between

shapes of parameterized surfaces. These geodesics are invariant to rigid motion, scaling and most importantly reparameterization of individual surfaces. The novelty lies in defining a Riemannian metric directly on the quotient (shape) space, rather than inheriting it from pre-shape space, and in using it to formulate a path energy that measures only the normal components of velocities along the path. The geodesic computation is based on a path-straightening technique that iteratively corrects paths between surfaces until geodesics are achieved. We have presented some examples of geodesics between surfaces in shape spaces and utilized the distances between surfaces for classification of some 3D shapes. However, the computational costs of our programs are deemed high and convergence should be accelerated in order to be able to apply this framework in realistic practical scenarios such as, for instance, human body action recognition.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for valuable comments that helped improve the paper significantly. The spherical parameterizations of the surfaces used in this paper were implemented by H. Laga. The authors would like to thank S. Kurtek for his help, and J.-C. Alvarez-Paiva and M. Bauer for fruitful discussions. This work was supported in part by the Labex CEMPI (ANR-11-LABX-0007-01) and was made possible by a visit of the first author to Telecom-Lille CRISTAL financed by the CNRS. During the reviewing process, the first author enjoyed excellent working conditions at the Pauli Institute, Vienna, Austria, and the first and last author benefit of the program on Shape Analysis held in 2015 at the ESI, Wien, Austria.

REFERENCES

- [1] T. Windheuser, U. Schlickewei, F. R. Schmidt, and D. Cremers, "Geometrically consistent elastic matching of 3d shapes: A linear programming solution," in *Proceedings of the 2011 International Conference on Computer Vision*, ser. ICCV '11, 2011, pp. 2134–2141.
- [2] M. Kilian, N. J. Mitra, and H. Pottmann, "Geometric modeling in shape space," *ACM Trans. Graphics*, vol. 26, no. 3, 2007, Proc. SIGGRAPH.
- [3] B. Heeren, M. Rumpf, M. Wardetzky, and B. Wirth, "Time-discrete geodesics in the space of shells," *Comp. Graph. Forum*, vol. 31, no. 5, pp. 1755–1764, Aug. 2012.
- [4] L. Younes, "Computable elastic distances between shapes," *SIAM J. of Applied Math*, pp. 565–586, 1998.
- [5] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn, "Shape analysis of elastic curves in euclidean spaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 7, pp. 1415–1428, 2011.
- [6] S. Kurtek, E. Klassen, Z. Ding, and A. Srivastava, "A novel Riemannian framework for shape analysis of 3d objects," in *CVPR*, 2010, pp. 1625–1632.
- [7] S. Kurtek, E. Klassen, J. C. Gore, Z. Ding, and A. Srivastava, "Elastic geodesic paths in shape space of parameterized surfaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 9, pp. 1717–1730, 2012.
- [8] N. Litke, M. Droske, M. Rumpf, and P. Schröder, "An image processing approach to surface matching," in *Proceedings of the Third Eurographics Symposium on Geometry Processing*, ser. SGP '05, 2005.
- [9] I. H. Jermyn, S. Kurtek, E. Klassen, and A. Srivastava, "Elastic shape matching of parameterized surfaces using square root normal fields," in *ECCV (5)*, 2012, pp. 804–817.
- [10] P. W. Michor and D. Mumford, "Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms," *Doc. Math.*, vol. 10, pp. 217–245, 2005.
- [11] M. Bauer, P. Harms, and P. W. Michor, "Curvature weighted metrics on shape space of hypersurfaces in n-space," *Differential Geom. Appl.*, vol. 30, pp. 33–41, 2012.
- [12] M. Bauer, P. Harms, and P. W. Michor, "Almost local metrics on shape space of hypersurfaces in n-space," *SIAM J. Img. Sci.*, vol. 5, no. 1, pp. 244–310, Mar. 2012.
- [13] M. Bauer, P. Harms, and P. W. Michor, "Sobolev metrics on shape space of surfaces," *Journal of Geometric Mechanics*, vol. 3, no. 4, pp. 389–438, 2011.
- [14] M. Fuchs, B. Jttler, O. Scherzer, and H. Yang, "Shape metrics based on elastic deformations," *J. Math. Imaging Vis.*, vol. 35, no. 1, pp. 86–102, 2009.
- [15] M. Bauer, M. Bruveris, and P. W. Michor, "Overview of the geometries of shape spaces and diffeomorphism groups," *Journal of Mathematical Imaging and Vision*, vol. 50, no. 1-2, pp. 60–97, 2014.
- [16] D. G. Ebin, "The manifold of Riemannian metrics," *Global Analysis (Proc. Sympos. Pure Math., Vol. XV, Berkeley, Calif., 1968)*, pp. 11–40, 1970.
- [17] D. S. Freed and D. Groisser, "The basic geometry of the manifold of Riemannian metrics and of its quotient by the diffeomorphism group," *Michigan Math. J.*, 1989.
- [18] O. Gil-Medrano and P. W. Michor, "The Riemannian manifold of all Riemannian metrics," *Quarterly J. Math. Oxford*, vol. 2, 1991.
- [19] B. Clarke, "The metric geometry of the manifold of Riemannian metrics over a closed manifold," *Calculus of Variations and Partial Differential Equations*, vol. 39:533–545, 2010.
- [20] M. Bauer, M. Bruveris, and P. W. Michor, "Constructing reparametrization invariant metrics on spaces of plane curves," *Differential Geometry and its Applications*, vol. 34, pp. 139–165, 2014.
- [21] S. Lahiri, D. Robinson, and E. Klassen, "Precise matching of PL curves in \mathbb{R}^N in the square root velocity framework," 2015.
- [22] E. Binz and H. R. Fischer, "The manifold of embeddings of a closed manifold," in *Proc. Differential geometric methods in theoretical physics, Clausthal 1978*, vol. Springer Lecture Notes in Physics 139, 1981.
- [23] P. W. Michor, *Manifolds of differentiable mappings*. Shiva Publishing Ltd., Nantwich, 1980.
- [24] M. P. do Carmo, *An Introduction to Differential Geometry with Applications to Elasticity*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1976.
- [25] W. Klingenberg, *Eine Vorlesung in Differentialgeometrie*. Springer Verlag, Berlin, 1973.
- [26] P. G. Ciarlet, *An Introduction to Differential Geometry with Applications to Elasticity*. Kluwer Academic Publishers, 2005, vol. 78-79, no. 1-3.
- [27] I. Ivanova-Karatopraklieva and I. K. Sabitov, "Surface deformation. i." *Journal of Mathematical Sciences*, vol. 70, no. 2, pp. 1685–1716, 1994.
- [28] J. Koenderink, *Solid Shape*. MIT Press, 1990.
- [29] R. Courant and D. Hilbert, *Methods of Mathematical Physics*. Wiley-Interscience, 1962, vol. Volume I.
- [30] Q. Xie, S. Kurtek, H. Le, and A. Srivastava, "Parallel transport of deformations in shape space of elastic surfaces," in *Proceedings of the 2013 International Conference on Computer Vision*, ser. ICCV '13, 2013.
- [31] A. Bronstein, M. Bronstein, and R. Kimmel, *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.
- [32] S. Kurtek, A. Srivastava, E. Klassen, and H. Laga, "Landmark-guided elastic shape analysis of spherically-parameterized surfaces," *Eurographics*, vol. 32, no. 2, 2013.



Alice Barbara Tumpach is an Associate Professor in Mathematics (University Lille 1, France) and member of the Laboratoire Painlevé (Lille 1/CNRS UMR 8524), since 2007. She received a Ph.D degree in Mathematics in 2005 at the Ecole Polytechnique, Palaiseau, France. She spent two years at the Ecole Polytechnique Fédérale de Lausanne as a Post-Doc.

Her research interests lie in the area of infinite-dimensional Geometry, Lie Groups and Functional Analysis.



Hassen Drira is an assistant Professor of Computer Science at Télécom Lille and member of the Laboratoire CRISTAL (UMR CNRS 9189) since September 2012. He obtained his Ph.D degree in Computer Science in 2011 from University of Lille 1 (France). His research interests are mainly focused on pattern recognition, statistical shape analysis. He has published several refereed journals and conference articles in these areas.



Mohamed Daoudi is a Professor of Computer Science at Telecom Lille and member of the Laboratoire CRISTAL (UMR CNRS 9189). He received his Ph.D. degree in Computer Engineering from the University of Lille 1 (France) in 1993 and Habilitation Diriger des Recherches from the University of Littoral (France) in 2000. His research interests include

pattern recognition, shape analysis and computer vision. He is a Fellow member of the IAPR.



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.