

A numerical framework for elastic surface matching, comparison, and interpolation

Martin Bauer · Nicolas Charon · Philipp Harms · Hsi-Wei Hsieh

Received: DD Month YEAR / Accepted: DD Month YEAR

Abstract Surface comparison and matching is a challenging problem in computer vision. While elastic Riemannian metrics provide meaningful shape distances and point correspondences via the geodesic boundary value problem, solving this problem numerically tends to be difficult. Square root normal fields (SRNF) considerably simplify the computation of certain distances between parametrized surfaces. Yet they leave open the issue of finding optimal reparametrizations, which induce corresponding distances between unparametrized surfaces. This issue has concentrated much effort in recent years and led to the development of several numerical frameworks. In this paper, we take an alternative approach which bypasses the direct estimation of reparametrizations: we relax the geodesic boundary constraint using an auxiliary parametrization-blind varifold fidelity metric. This reformulation has several notable benefits. By avoiding altogether the need for repara-

MB is supported by NSF grant 1912037 (collaborative research in connection with 1912030). NC and HH are supported by NSF grants 1819131 and 1945224.

M. Bauer
Florida State University
Department of Mathematics
E-mail: bauer@maths.fsu.edu

N. Charon
Johns Hopkins University
Department of Applied Mathematics and Statistics
E-mail: charon@cis.jhu.edu

P. Harms
University of Freiburg
Department of Mathematical Stochastics
E-mail: philipp.harms@stochastik.uni-freiburg.de

H-W. Hsieh
Johns Hopkins University
Department of Applied Mathematics and Statistics
E-mail: hhsieh@cis.jhu.edu

metrizations, it provides the flexibility to deal with simplicial meshes of arbitrary topologies and sampling patterns. Moreover, the problem lends itself to a coarse-to-fine multi-resolution implementation, which makes the algorithm scalable to large meshes. Furthermore, this approach extends readily to higher-order feature maps such as square root curvature fields and is also able to include surface textures in the matching problem. We demonstrate these advantages on several examples, synthetic and real.

Keywords elastic shape analysis, surfaces, square root normal fields, varifolds.

Contents

1	Introduction	1
2	Method	4
3	Elastic interpretation	8
4	Implementation	9
5	Numerical results	13
6	Discussion and conclusion	17
A	Notation	18
B	Formula for SRNF metrics	18
C	Approximation properties of SRNF distances	18
D	Formula for SRCF metrics	19

1 Introduction

The analysis of shape-and-image data is an important problem in computer vision. Such data arise naturally in many applications such as biomedicine (Pennec et al. 2019) or robotics (Turaga and Srivastava 2016) and is produced e.g. by medical imaging devices or time-of-flight cameras in cell phones or cars (Grzegorzek et al.

2013). While deep-learning techniques have led to significant break-throughs in image analysis, these developments have not yet been paralleled in shape analysis (cf. Geirhos et al. 2018). A major difficulty in the direct application of deep-learning methods to shapes is that common shape descriptors are non-unique and vary in structure and dimension. For example, one and the same surface can be represented by many different triangular meshes with varying connectivities and numbers of vertices. While differences in mesh parametrizations matter in numerical computations, they carry no statistical information about the shapes themselves and have to be quotiented out in meaningful statistical analyses.

Shape analysis provides a mathematical description of shape spaces as quotient spaces in the above sense, as well as a computational toolbox for statistics and machine learning thereon. Some classical textbooks are Dryden and Mardia (1998) and Kendall et al. (1999), and some more recent ones Younes (2010), Srivastava and Klassen (2016), and Pennec et al. (2019). In shape analysis, as in functional data analysis (Kokoszka and Reimherr 2017), the starting point is typically a space of functions, say from S^2 to \mathbb{R}^3 in the case of surfaces.

In numerical implementations, this function space is discretized, e.g. by triangular meshes, seen as piecewise linear \mathbb{R}^3 -valued functions on some fixed triangulation of S^2 . The shape space corresponding to this function space is determined by the following equivalence relation: two functions are considered as representations of the same shape if they differ only by a reparametrization, i.e., by a diffeomorphism on their domain. Additionally, rigid motions and scalings are sometimes factored out, as well (Kendall et al. 1999). Shape space is a nonlinear infinite-dimensional manifold (Cervera et al. 1991), even if the original function space is linear, and it is the natural configuration space for statistics and machine learning on shapes.

Riemannian geometry is ideally suited as a basis for statistics and machine learning on manifolds of shapes or otherwise: it allows the estimation of Fréchet means and higher-order statistical moments via the geodesic distance function (Dryden and Mardia 1998), provides local linearizations and geodesic principal components via the Riemannian logarithm (Srivastava and Klassen 2016), defines generative random models via stochastic geometric mechanics (Pennec et al. 2019), and supplies kernels for support vector machines via the geodesic distance and heat equation (Minh et al. 2016).

The intuition for the Riemannian setting is that two shapes are considered to be similar when they differ only by a small deformation, as measured by a Riemannian metric.

The Riemannian metrics studied in this paper are often called elastic (Younes 1998; Srivastava et al. 2011; Kurtek et al. 2012; Jermyn et al. 2012, 2017) despite the analogy to elastic stretching and bending energies being only loose (Rumpf and Wardetzky 2014; Rumpf and Wirth 2015a). Alternatively, these metrics are also called inner metrics (Bauer et al. 2011) to distinguish them from outer metrics in diffeomorphic shape analysis (Younes 2010). All elastic metrics are invariant to reparametrizations, and this is needed for the metrics to be well-defined on the quotient space of shapes (Bauer et al. 2016).

Reparametrization-invariance of the metric implies that even linear function spaces have interesting geometries, which are full of surprises. For example, the geodesic distance of the simplest reparametrization-invariant Riemannian metric vanishes on spaces of curves, surfaces, and diffeomorphisms, as discovered by Michor and Mumford (2005, 2006). This degeneracy is a purely infinite-dimensional phenomenon, which led to a systematic study of higher-order Sobolev metrics on mapping spaces. It turns out that metrics of order one and higher have non-vanishing geodesic distance (Bauer et al. 2011), and metrics of order two and higher are complete on spaces of curves with the same Sobolev regularity (Bruveris et al. 2014). The metrics considered here are of similar type but lack a zero-order term, which means that they are blind to translations and, beyond dimension one, some further non-trivial deformations (Klassen and Michor 2020).

In data-analytic applications, what is needed foremost are efficient implementations of the geodesic initial and boundary value problems on shape space (see Section 3.1 for definitions and Srivastava and Klassen (2016) for applications). The theoretical understanding of these problems is still incomplete. The initial value problem is locally well posed for a wide class of Sobolev metrics of order one and higher (Bauer et al. 2011), even of fractional order (Bauer et al. 2020). Moreover, the initial and boundary value problems are globally well-posed for certain second-order Sobolev metrics on spaces of curves (Bruveris 2015), but this has not yet been generalized to surfaces. Algorithmically, the initial value problem can be solved by a time-discretization of the geodesic equation (Bauer et al. 2011) or using discrete geodesic calculus (Rumpf and Wirth 2015b). The boundary value problem on functions (seen as parameterized surface) is closely related to mesh interpolation (Kilian et al. 2007; Fröhlich and Botsch 2011) and can be solved by the direct method of variational calculus or, for certain elastic metrics, by exploiting isometries to simpler spaces. For the boundary value problem on shape spaces, one has to solve in addition the optimal

correspondence problem, i.e., one has to factor out the action of the group of reparametrizations.

Searching for optimal reparametrizations is a challenging task. Previous approaches relied on a discretization of the reparametrization group and an implementation of the group action on the discretized function space (Srivastava and Klassen 2016; Jermyn et al. 2017; Su et al. 2020b). Unfortunately, the implementation of the group action requires the initial and target shapes to be defined on the same pre-specified domain with pre-specified triangulation. As this is typically not the case, one first has to solve the parametrization problem (Sheffer et al. 2007), which is of comparable difficulty to the geodesic boundary problem itself. An additional problem is that the reparametrization group can be discretized well for circular domains using monotone correspondences (Bernal et al. 2016) and for spherical domains using spherical harmonics (Jermyn et al. 2012), but adaptations to more general domains are difficult. Furthermore, compressions of surface patches to points and decompressions of points to surface patches are non-trivial to implement but occur in optimal reparametrizations (Lahiri et al. 2015). Some of these problems can be avoided by minimizing a gauge-invariant energy (Tumpach et al. 2015; Tumpach 2016), but this cannot be used to compute optimal point correspondences, which is one of our main goals.

The main contribution of this paper is a reformulation of the geodesic boundary value problem, which circumvents the above-mentioned problems. This reformulation draws on two lines of work: square root normal fields and varifold distances. Square root normal fields are isometric transformations from spaces of curves or surfaces to simpler manifolds. They were introduced by Kurtek et al. (2012) and Jermyn et al. (2012) and have their origin in similar transformations for curves, which were discovered by Younes (1998), Sundaramoorthy et al. (2011), and Srivastava et al. (2011). Square root normal fields map surfaces isometrically into a Hilbert space of square-integrable vector-valued half-densities. The Hilbert distance between square root normal fields is a first-order approximation of a Riemannian distance and can be computed efficiently by evaluating an integral. In contrast to general Riemannian distances, no optimization or time discretization is involved. Thus, square root normal fields provide computable (approximate) Riemannian distances with an elastic interpretation.

Varifold distances have their origin in geometric measure theory (Federer 1969; Almgren 1966) and were introduced into the context of diffeomorphic shape analysis by Charon and Trouvé (2013), building on previous work on current distances by Vaillant and Glaunès

(2005) and Glaunès et al. (2008). Their main advantages are that they are fully blind to reparametrizations and can be implemented efficiently on parallel computing architectures, as shown by Kaltenmark et al. (2017) and Charon et al. (2020). Despite directly providing distances on shape spaces, the main limitation is that those distances are not the result of a Riemannian metric: in particular, there is no interpretable notion of geodesics between two shapes in this setting. For this reason, they have to be combined with other shape distances, either on diffeomorphism groups as in diffeomorphic matching (Charon and Trouvé 2013) or on shape spaces as in elastic matching. The latter approach was used for Sobolev metrics on curves by Bauer et al. (2017, 2019a) and for square root normal distances on curves in the recent conference paper by Bauer et al. (2019b), which is a predecessor of the present paper. In the present article, the focus is on surfaces, and we add multi-resolution matching, square root normal field inversion, geodesic interpolation, square root curvature fields, surface noise, textured surfaces, non-homotopic surfaces, and Fréchet means. Moreover, we provide a stable open-source implementation in **Python**.

We combine square root normal fields and varifold distances in a new elastic shape matching algorithm, i.e., a new algorithm for solving the geodesic boundary value problem of elastic metrics on shape space. The idea is to use varifold distances for relaxing the boundary constraint in the elastic matching problem and to use square root normal fields for boosting the computation of Riemannian distances. This has the following advantages.

- **Speed:** Square root normal fields allow one to bypass the time discretization of the geodesic equation, and varifold distances circumvent the costly discretization of reparametrizations.
- **Applicability and flexibility:** The algorithm can be applied directly to simplicial meshes without having to solve the parametrization problem first. It can handle different mesh structures and even different topologies, which is important in the presence of topological noise or for matching shapes with missing parts. Moreover, texture information can be used as in the **fshape** framework of Charon and Trouvé (2014) and Charlier et al. (2017).
- **Correctness:** The point correspondences found by our algorithm may exhibit compression of surface patches to points and decompression of points to surface patches. Such compressions and decompressions may indeed occur in optimal point correspondences (at least for curves; see Lahiri et al. 2015) and are difficult to model in explicit discretizations of diffeomorphism groups.

- Robustness and scalability: The multi-resolution version of our algorithm robustly identifies optimal point correspondences, even for high-resolution surfaces. It scales well to high dimensions thanks to the parallelizability of varifold computations and the linear computational complexity of square root normal fields.

All algorithms are publicly available¹ as a Python package **SRNFmatch**.

1.1 Structure of the paper

The paper is structured as follows. Section 2 describes our new problem formulation for the computation of square root normal field distances between surfaces. Section 3 connects this approach to solving the geodesic boundary value problem of a specific elastic shape metric. Section 4 presents the discretization of the algorithm on triangular meshes as well as a multi-resolution scheme to tackle it more efficiently. Section 5 demonstrates the usefulness of our algorithm on real-world and artificial test data.

2 Method

This section describes the proposed new algorithm for the computation of square root normal distances between shapes, using varifold distances as an auxiliary tool. An elastic interpretation is given subsequently in Section 3. The notation is summarized in Appendix A.

2.1 Shape spaces

This section defines spaces of parametrized and unparametrized surfaces. The latter ones are also called shapes and correspond to equivalence classes of parametrized surfaces modulo reparametrizations. Additionally, following Kendall et al. (1999), one could factor out translations and rotations. However, as translation and rotation groups are only finite-dimensional, we focus on the more difficult task of quotienting out the infinite-dimensional reparametrization group.

To make this precise, we define the space of parametrized surfaces as the set \mathcal{I} of all oriented immersions q of a 2-dimensional compact manifold M (possibly with boundary) into \mathbb{R}^3 . The set \mathcal{I} is an open subset of the Fréchet space $C^\infty(M, \mathbb{R}^3)$. The diffeomorphism group of M is denoted by \mathcal{D} and acts on \mathcal{I} by reparametrization, i.e., by composition from the right.

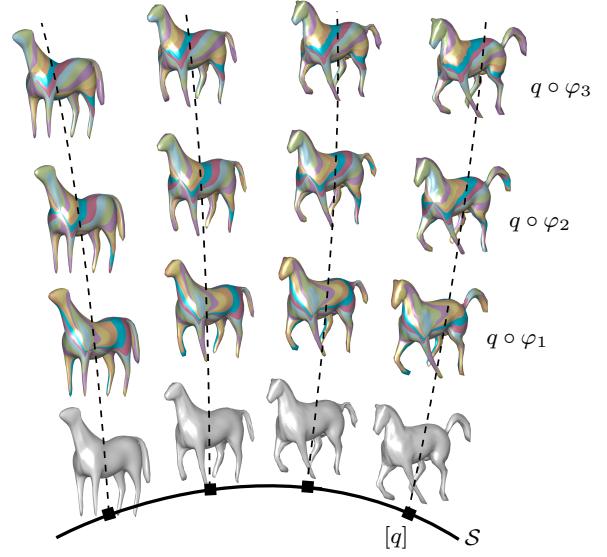


Fig. 1 Shape space as a quotient of the space of parametrized surfaces. Each shape $[q]$ is an equivalence class of surfaces $q \circ \varphi$, which differ only by a reparametrization φ . For example, the shape of a horse (gray) is an equivalence class (dashed line) of surfaces with different mesh parametrizations (color-coded). Elastic shape matching requires the computation of the minimal distance between surfaces of one equivalence class to surfaces of another equivalence class.

The quotient space $\mathcal{S} = \mathcal{I}/\mathcal{D}$, which is called shape space, consists of unparametrized surfaces, i.e., equivalence classes $[q] = \{q \circ \varphi; \varphi \in \mathcal{D}\}$. This is illustrated in Figure 1: each equivalence class is a set of parametrized surfaces which correspond to one and the same shape because they differ only by a reparametrizations. On a discrete level, \mathcal{I} corresponds to triangular meshes with fixed connectivity, and reparametrizations correspond to remeshings.

Unfortunately, the quotient space \mathcal{S} has no straightforward discretization, and this constitutes one of the main difficulties in numerical shape analysis.

2.2 Square root normal fields

Square root normal fields are part of a family of feature maps, which can be used to define shape distances related to elastic metrics, as we explain in Section 3.2. The first instance of such a feature map, which was discovered by Younes et al. (2008) and Sundaramoorthy et al. (2011), is the map from planar curves q to the complex square root $\sqrt{q'}$ of their velocity, seen as an element of an infinite-dimensional Stiefel manifold. Similarly, Srivastava et al. (2011) considered the map from planar curves q to their square root velocity field $q'|q'|^{-1/2}$, seen as an element of a Hilbert space. Subse-

¹ https://github.com/SRNFMATCH/SRNFMATCH_code

quently, Kurtek et al. (2012) and Jermyn et al. (2012) found an extension to higher dimension, namely the map from oriented hyper-surfaces q to their square root normal field (SRNF) N_q , which is the unit normal field n_q multiplied by the square root of the Riemannian area form A_q :

$$N_q := n_q A_q^{1/2}.$$

For curves, the SRNF coincides up to a rotation of 90 degrees with the square root velocity field. For two-dimensional surfaces, it is given in coordinates $(u, v) \in \mathbb{R}^2$ as

$$N_q = (q_u \times q_v) |q_u \times q_v|^{-1/2}, \quad (1)$$

where the subscripts denote partial derivatives, \times denotes the cross product on \mathbb{R}^3 , and $|\cdot|$ denotes the norm on \mathbb{R}^3 . Geometrically, the SRNF belongs to the space of square-integrable \mathbb{R}^3 -valued half-densities, which is a Hilbert space with the L^2 scalar product (see Michor 2008, Section 10.4 for further details). Half-densities can be identified with real-valued functions by fixing a reference half-density. This shall be done implicitly in all numerical discretizations. However, it should be kept in mind that half-densities transform differently than functions under reparametrizations: for any $\varphi \in \mathcal{D}$, in any coordinate system on M ,

$$N_{q \circ \varphi} = (N_q) \circ \varphi \det(D\varphi)^{1/2}, \quad (2)$$

where $D\varphi$ denotes the differential (or Jacobian) of the diffeomorphism φ .

The SRNF distance between parametrized surfaces q_0 and q_1 is defined as the L^2 distance between the associated SRNFs:

$$\text{dist}_N(q_0, q_1) := \int_M |N_{q_0} - N_{q_1}|^2. \quad (3)$$

Note that the right-hand side is well-defined because the square of a half-density is a density. The SRNF distance is sometimes called chordal distance to point out that the SRNF map is non-surjective and that consequently the SRNF distance differs from the intrinsic distance on its range. The SRNF distance has two advantages: it can be computed easily, as described in Section 4, and it is \mathcal{D} -invariant, i.e., for any $\varphi \in \mathcal{D}$,

$$\text{dist}_N(q_0 \circ \varphi, q_1 \circ \varphi) = \text{dist}_N(q_0, q_1). \quad (4)$$

This follows from (2) and the change-of-variable formula for integrals. However, this distance is not yet a distance between shapes, and it remains to quotient out reparametrizations. Accordingly, the SRNF distance between shapes $[q_0], [q_1] \in \mathcal{S}$ is defined as the minimal

SRNF distance between all reparametrizations of q_0 and q_1 , i.e.,

$$\begin{aligned} \text{dist}_N([q_0], [q_1]) &:= \inf_{\varphi_0, \varphi_1 \in \mathcal{D}} \text{dist}_N(q_0 \circ \varphi_0, q_1 \circ \varphi_1) \\ &= \inf_{\varphi \in \mathcal{D}} \text{dist}_N(q_0, q_1 \circ \varphi), \end{aligned} \quad (5)$$

where the second equation follows from the \mathcal{D} -invariance (4). Similar distances can be defined for more general \mathcal{D} -equivariant feature maps, as discussed in Section 3.3. An important caveat is that the SRNF distance is only a pseudo-distance because the SRNF map is non-injective, even after factoring out translations; see Klassen and Michor (2020). We will not dwell upon this issue and freely speak of distances even when they are only pseudo-distances. The main difficulty in (5) is the search for an optimal reparametrization $\varphi \in \mathcal{D}$. We address this problem using varifold distances, which are described next.

2.3 Varifold distances

Geometric measure theory provides several embeddings of shape spaces into some Banach spaces of distributions, including currents (Vaillant and Glaunès 2005; Glaunès et al. 2008), varifolds (Charon and Trouvé 2013), and normal cycles (Roussillon and Glaunès 2016), see (Charon et al. 2020) for a recent review. In this paper, we adopt the framework of Kaltenmark et al. (2017). The varifold μ_q associated to $q \in \mathcal{I}$ is the image measure $(q, n_q)_* A_q$ on $\mathbb{R}^3 \times S^2$, where n_q is the unit normal field of q , and A_q is the area form (also known as surface measure) of q . In other words, for any Borel set $B \subset \mathbb{R}^3 \times S^2$, $\mu_q(B)$ is the area of all $m \in M$ such that $(q(m), n_q(m))$ belongs to B . Importantly, this measure representation does not depend on the parametrization of q and thus provides an embedding of shape space \mathcal{S} into the space of positive measures of $\mathbb{R}^3 \times S^2$. Then, given a norm $\|\cdot\|$ on the space of measures, one defines the varifold distance between $[q_0], [q_1] \in \mathcal{S}$ as

$$\text{dist}_V(q_0, q_1) := \|\mu_{q_0} - \mu_{q_1}\|. \quad (6)$$

While there are many possible distances that one can introduce on spaces of measures, norms defined from positive definite kernels on $\mathbb{R}^3 \times S^2$ have been shown to lead to particularly advantageous expressions for numerical computations. Specifically, following the setting of Kaltenmark et al. (2017), we consider the class of norms $\|\cdot\|_{V^*}$, where V is a reproducing kernel Hilbert space of functions on $\mathbb{R}^3 \times S^2$, whose kernel is of the form

$$k(x_1, n_1, x_2, n_2) := \rho(|x_1 - x_2|) \gamma(n_1 \cdot n_2), \quad (7)$$

where ρ and γ are two functions defining a radial kernel on \mathbb{R}^3 and a zonal kernel on S^2 , respectively. Then the scalar product between varifolds μ_{q_0} and μ_{q_1} can be shown to be:

$$\begin{aligned} \langle \mu_{q_0}, \mu_{q_1} \rangle_{V^*} &= \iint_{M \times M} \rho(|q_0(u_0, v_0) - q_1(u_1, v_1)|) \\ &\quad \gamma(n_{q_0}(u_0, v_0) \cdot n_{q_1}(u_1, v_1)) A_{q_0}(u_0, v_0) A_{q_1}(u_1, v_1) \end{aligned} \quad (8)$$

Accordingly, the varifold distance between any two triangular surfaces q_0 and q_1 can be computed as

$$\text{dist}_V(q_0, q_1)^2 = \|\mu_{q_0}\|_{V^*}^2 - 2\langle \mu_{q_0}, \mu_{q_1} \rangle_{V^*} + \|\mu_{q_1}\|_{V^*}^2.$$

As shall be seen in Section 4, the previous expression has a natural and simple discrete equivalent for triangular meshes.

Varifold distances are equivariant to the action of rigid motions: for any $R \in SO(3)$ and $h \in \mathbb{R}^3$,

$$\text{dist}_V(Rq_0 + h, Rq_1 + h) = \text{dist}_V(q_0, q_1).$$

More importantly, they are $(\mathcal{D} \times \mathcal{D})$ -invariant: for any $\varphi_0, \varphi_1 \in \mathcal{D}$,

$$\text{dist}_V(q_0 \circ \varphi_0, q_1 \circ \varphi_1) = \text{dist}_V(q_0, q_1).$$

There is an important difference between the \mathcal{D} -invariance of the SRNF distance and the $(\mathcal{D} \times \mathcal{D})$ -invariance of varifold distances: a \mathcal{D} -invariant distance vanishes at $q_0, q_1 \in \mathcal{I}$ if and only if $q_0 = q_1$, whereas a $(\mathcal{D} \times \mathcal{D})$ -invariant distance vanishes at $q_0, q_1 \in \mathcal{I}$ if and only if $q_0 = q_1 \circ \varphi$ for some $\varphi \in \mathcal{D}$. Varifold distance are rather ideally suited for enforcing the latter condition, i.e., membership of the same \mathcal{D} -orbit in \mathcal{I} . This makes them highly useful in shape analysis, as described next. However, they are not tied to a Riemannian metric and thus are not associated to geodesics between shapes.

Another interesting feature of the varifold framework is the possibility to incorporate additional texture information in the metric, as detailed in (Charlier et al. 2017). Assume that the given immersed shape q carries a scalar texture map $\zeta: M \rightarrow \mathbb{R}$. Then the couple (q, ζ) can be represented as a functional varifold, i.e., as the image measure

$$\mu_{q, \zeta} := (q, n_q, \zeta)_* A_q$$

on $\mathbb{R}^3 \times S^2 \times \mathbb{R}$. The previous kernel metrics can then be readily extended to this new situation using the following new kernel:

$$k(x_1, n_1, \zeta_1, x_2, n_2, \zeta_2) := \rho(|x_1 - x_2|) \gamma(n_1 \cdot n_2) \tau(|\zeta_1 - \zeta_2|)$$

Here, τ defines a positive-definite kernel on \mathbb{R} , which adds to (7) a notion of proximity between texture values. The corresponding functional varifold distance is

defined similarly to (8) and (6). It provides a distance between textured surfaces (q_0, ζ_0) and (q_1, ζ_1) , which can be used to enforce the joint constraint of matching geometries $q_0 = q_1 \circ \varphi$ and matching textures $\zeta_0 = \zeta_1 \circ \varphi$, up to reparametrization by some $\varphi \in \mathcal{D}$.

2.4 Combining SRNF and varifold distances

From a numerical perspective, the main difficulty in the computation (5) of SRNF distances between shapes is the search for an optimal reparametrization $\varphi \in \mathcal{D}$ in the diffeomorphism group of M . For curves, reparametrizations of $M = S^1$ are monotone correspondences, and the optimal correspondence can be found using dynamic programming, following Frenkel and Basri (2003) and Sebastian et al. (2003). For surfaces, diffeomorphisms of $M = S^2$ can be represented using spherical harmonics, and the optimal harmonic representation can be found via gradient descent (Srivastava et al. 2011; Kurtek et al. 2012; Jermyn et al. 2012). To implement the action of diffeomorphisms on shapes, these algorithms presuppose that the shapes are already parametrized consistently, i.e., defined on the same mesh. Otherwise, one has to solve the challenging parametrization problem first, which is of comparable difficulty to the search for an optimal reparametrization. An additional disadvantage is that these methods do not generalize easily to domains other than S^1 or S^2 .

We circumvent these problems by exploiting the $(\mathcal{D} \times \mathcal{D})$ -invariance of varifold distances. Previously, Charon and Trouvé (2013) and Kaltenmark et al. (2017) applied a similar strategy to diffeomorphic matching. Moreover, Bauer et al. (2019a) used varifold distances for the first time in the context of elastic matching, albeit only for curves. The starting point is the observation that problem (5) is equivalent to the constrained optimization problem

$$\min_{\tilde{q}_1 \in \mathcal{I}} \text{dist}_N(q_0, \tilde{q}_1) \quad \text{subject to} \quad \text{dist}_V(\tilde{q}_1, q_1) = 0.$$

The optimization over $\varphi \in \mathcal{D}$ in (5) has now been replaced by an optimization over $\tilde{q}_1 \in \mathcal{I}$ subject to a varifold constraint, which ensures that \tilde{q}_1 belongs to the \mathcal{D} -orbit of q_1 , as required. Relaxation of the constraint using a (large) penalty parameter λ yields

$$\min_{\tilde{q}_1 \in \mathcal{I}} \text{dist}_N(q_0, \tilde{q}_1) + \lambda \text{dist}_V(\tilde{q}_1, q_1). \quad (9)$$

At this point one may wonder whether both of the two distances in (9), namely the varifold and SRNF distance, are really needed. This is indeed the case, as the formulation exploits—and requires—both the $(\mathcal{D} \times \mathcal{D})$ -invariance of the varifold distance and the \mathcal{D} -invariance

of the SRNF distance. In other words, the SRNF distance alone cannot explore different parametrizations, while the varifold distance alone cannot distinguish between good and bad parametrizations and does not define geodesics between shapes.

2.5 Symmetric elastic-varifold matching

In the category of smooth surfaces and diffeomorphisms, the asymmetric matching problem (9) is equivalent to the following symmetrized matching problem:

$$\min_{\tilde{q}_0, \tilde{q}_1 \in \mathcal{I}} \lambda \text{dist}_V(q_0, \tilde{q}_0) + \text{dist}_N(\tilde{q}_0, \tilde{q}_1) + \lambda \text{dist}_V(\tilde{q}_1, q_1). \quad (10)$$

However, for non-smooth surfaces and diffeomorphisms the two problems are different, and the symmetric version offers significant advantages. The symmetric matching problem (10) is more general, as it is well-defined for arbitrary varifolds q_0 and q_1 . As an aside, it should be kept in mind that the matching energy (10) depends on the choice of M . This choice can be interpreted as a topological prior and corresponds on a discrete level to a choice of triangular mesh for the surfaces \tilde{q}_0 and \tilde{q}_1 . A further advantage is that the symmetric matching problem allows for reparametrizations of both surfaces q_0 and q_1 . In the case of curves this flexibility is needed to represent optimal reparametrizations, which may compress a patch to a point or conversely decompress a point to a patch (Lahiri et al. 2015). Similar compressions and decompressions are to be expected in the case of surfaces and cannot be modeled in the asymmetric matching problem, where only one of the surfaces is reparametrized.

2.6 Comparison to alternative methods

Compared to alternative methods of surface comparison and matching, program (10) has the following advantages:

- The reparametrization group \mathcal{D} does not need to be discretized, and its action on \mathcal{I} does not need to be implemented. This allows one to work with simplicial meshes without having to solve the parametrization problem first, as for instance in Srivastava and Klassen (2016) and Jermyn et al. (2017). Moreover, it easily generalizes to domains M more general than S^1 or S^2 , which is not without difficulties when dynamic programming or spherical harmonic are used, as in e.g. Bernal et al. (2016) or Jermyn et al. (2012).

- The varifold matching term provides the flexibility for handling topological noise, matching shapes with missing parts, and incorporating texture information, similarly to the `fshape` framework (Charon and Trouv   2014; Charlier et al. 2017).
- The computations are faster than for general Riemannian distances (see Section 3), where an additional time discretization is needed.

These advantages are also demonstrated by the numerical examples in Section 5.

2.7 Applications

Program (10) is the basis for several important tasks in shape analysis:

- It implements the geodesic boundary value problem for an elastic shape metric, which is described in further detail in Section 3. At intermediate time points the geodesic can be obtained from the linear interpolation of the SRNFs by approximate inversion of the SRNF map $q \mapsto N_q$, as described in Section 4.8.
- The geodesic boundary value problem is the basis for machine learning and statistics on manifolds. It provides geodesic distances, which are used for agglomerative clustering, support vector machines, Fr  chet means, and higher-order moments; cf. Srivastava and Klassen (2016). Moreover, it implements the Riemannian logarithm, which is used in geodesic principal component analysis and statistics on manifolds; cf. Bhattacharya and Bhattacharya (2012) and Pennec et al. (2019). The Riemannian logarithm at the reparametrized surface \tilde{q}_0 is the initial velocity of the above-mentioned horizontal geodesic. Moreover, the Riemannian logarithm at the original surface q_0 can be obtained from the asymmetric version (9) of the program.
- The program implements the matching problem, i.e., the optimizers \tilde{q}_0 and \tilde{q}_1 describe a point correspondence between the given shapes q_0 and q_1 .
- The program implements the parametrization problem, i.e., the optimizers \tilde{q}_0 and \tilde{q}_1 parametrize the given shapes q_0 and q_1 as piece-wise linear functions on a given mesh. As an aside, this parametrization is not determined uniquely by (10) and could therefore be selected freely using some additional criteria of angle, distance, or area preservation, which have been developed in the literature on mesh parametrization; see e.g. Floater and Hormann (2005) or Sheffer et al. (2007).

3 Elastic interpretation

The surface matching approach presented in the previous section approximately solves the geodesic boundary value problem of a certain first-order elastic metric on the shape space of surfaces, as we explain next.

3.1 Elastic shape analysis

To emphasize this connection, we start with some background on elastic metrics, referring to the surveys by Bauer et al. (2014) or Jermyn et al. (2017) for further details and references. Elastic shape analysis operates in a Riemannian framework, where infinitesimal shape deformations are measured by a Riemannian metric, which is often related to an elastic (or plastic) deformation energy. Given a Riemannian metric G on \mathcal{I} , one defines the Riemannian distance between two shapes $q_0, q_1 \in \mathcal{I}$ as

$$\text{dist}_G(q_0, q_1)^2 := \inf_{\substack{q \in C^\infty([0,1], \mathcal{I}) \\ q(0)=q_0, q(1)=q_1}} \int_0^1 G_q(\partial_t q, \partial_t q) dt. \quad (11)$$

The optimizers in (11) are constant-speed geodesics in \mathcal{I} . If G is \mathcal{D} -invariant, as shall be assumed throughout, one obtains also a distance between shapes $[q_0], [q_1] \in \mathcal{S}$ by setting

$$\begin{aligned} \text{dist}_G([q_0], [q_1]) &:= \inf_{\varphi_0, \varphi_1 \in \mathcal{D}} \text{dist}_G(q_0 \circ \varphi_0, q_1 \circ \varphi_1) \\ &= \inf_{\varphi \in \mathcal{D}} \text{dist}_G(q_0, q_1 \circ \varphi), \end{aligned} \quad (12)$$

where the second equation follows from the \mathcal{D} -invariance of the Riemannian metric G . This is the same quotient construction as for SRNF distances (5), and the optimization over reparametrizations can again be avoided using varifold distances similarly to (10). However, in contrast to SRNF distances and elastic deformation energies in physics (Rumpf and Wardetzky 2014; Rumpf and Wirth 2015a), the Riemannian distance (12) retains a dynamical interpretation, as its optimizers are constant-speed geodesics in \mathcal{S} .

For low-order Sobolev metrics G , the distances on \mathcal{I} and \mathcal{S} do not separate points, as discovered by Michor and Mumford (2007). Fortunately, this degeneracy disappears for metrics of order one and higher (Bauer et al. 2011). For simplicity, we will speak of Riemannian metrics and their associated distances even when they are only pseudo-metrics and pseudo-distances.

From an applied perspective, efficient numerical implementations of the optimization problems (11) and (12) are crucial, as they are the algorithmic basis for the

computation of point correspondences, Fréchet means, geodesic principal components, parallel transport, and so on; see e.g. the survey of Bauer et al. (2014) or the book of Srivastava and Klassen (2016). For general elastic metrics, problem (11) can be implemented using path straightening methods. This involves a time discretization and is computationally costly. However, for certain elastic metrics, which are related to SRNFs, this time discretization can be avoided, as described next.

3.2 Square root normal metrics

Recall that the SRNF is a map from $q \in \mathcal{I}$ to the Hilbert space of square-integrable \mathbb{R}^3 -valued half-densities. The SRNF metric on \mathcal{I} is the pull-back of the Hilbert scalar product along this map. Thus, it is defined for any tangent vector $h \in T_q \mathcal{I}$ as

$$G_q(h, h) := \int_M |D_{(q,h)} N_q|^2, \quad (13)$$

where $D_{(q,h)} N_q$ denotes the directional derivative of N_q at q in the direction h . Note that the right-hand side in (13) is well-defined since the square of a half-density is a density.

The SRNF metric belongs to the class of first order Sobolev metrics, which have been studied in great detail by Michor and Mumford (2007), Mennucci et al. (2008), Bauer et al. (2011), and many others. An explicit formula for it is established in Appendix B:

$$G_q(h, h) = \int_M \left(|(\nabla h)^\perp|^2 + \frac{1}{4} \text{Tr}((\nabla h)^\top)^2 \right) A_q. \quad (14)$$

Here ∇ is the coordinate-wise derivative of \mathbb{R}^3 -valued functions, \perp and \top are the normal and tangential projections satisfying $\perp + Tq \circ \top = \text{Id}_{\mathbb{R}^3}$, the norm $|\cdot|$ is computed with respect to the pull-back cometric, and Tr denotes a trace; see Appendix A for further notation.

While computing geodesic distances of general elastic metrics can be quite cumbersome, this is much simpler for SRNF metrics thanks to the approximation

$$\text{dist}_G(q_0, q_1) \approx \text{dist}_N(q_0, q_1), \quad (15)$$

where dist_G is the Riemannian distance (11) of the SRNF metric (13), and dist_N is the SRNF distance (3). The approximation is exact whenever the straight line between N_{q_0} and N_{q_1} is contained in the range of the SRNF map $q \mapsto N_q$. The reason is that the SRNF map is a Riemannian isometry, by construction, and that geodesics in the Hilbert space of SRNFs are straight lines. In general, the approximation is exact up to first order for q_0 close to q_1 ; see Appendix C for a proof. Some generalizations to higher-order metrics are discussed next.

3.3 Generalized square root normal fields

The theoretical and algorithmic framework of the previous sections can be extended using more general feature maps than SRNFs. For example, Jermyn et al. (2012) introduced the Gauss map $q \mapsto (g_q, n_q)$ as a feature map, where g_q is the pull-back metric of q . This feature map is injective up to isometries of \mathbb{R}^3 by a classical result of Abe and Erbacher (1975), and the geodesic distance on the image space has recently been computed by Su et al. (2020a). Alternatively, Bauer et al. (2011) used the map $q \mapsto dq$ as a feature map, where dq is the derivative of q , seen as an \mathbb{R}^3 -valued one-form on M . For simplicity we restrict ourselves to feature maps which transform under reparametrizations similarly to half-densities. As explained below, this automatically guarantees \mathcal{D} -invariance of the L^2 distance between features, which is particularly easy to compute.

The general setup is as follows. One considers \mathcal{D} -equivariant maps Φ from \mathcal{I} to the Hilbert space of square-integrable sections of $E \otimes \text{Vol}^{1/2}$, where E is a tensor bundle over M with \mathcal{D} -invariant fiber metric $\langle \cdot, \cdot \rangle_E$ and corresponding fiber norm $|\cdot|_E$, and where $\text{Vol}^{1/2}$ is the half-density bundle over M . The corresponding distance, which generalizes the SRNF distance (3), is defined as

$$\text{dist}_\Phi(q_0, q_1)^2 := \int_M |\Phi_{q_0} - \Phi_{q_1}|_E^2. \quad (16)$$

Note that the integral on the right-hand side is well-defined because the square of a half-density is a density. Importantly, the distance can be computed easily from fiber-wise distances of the bundle E thanks the L^2 structure (cf. Ebin 1970). In particular, if $\langle \cdot, \cdot \rangle_E$ is flat on each fiber of E , then the computation reduces to the evaluation of an integral. Moreover, thanks to the assumption that Φ transforms as a half-density, the distance (16) is \mathcal{D} -invariant, i.e., for all $\varphi \in \mathcal{D}$,

$$\text{dist}_\Phi(q_0 \circ \varphi, q_1 \circ \varphi) = \text{dist}_\Phi(q_0, q_1).$$

The corresponding \mathcal{D} -invariant elastic metric, which generalizes the SRNF metric (13), is given by

$$G_q(h, h) := \int_M |D_{(q, h)} \Phi_q|_E^2, \quad (17)$$

where $h \in T_q \mathcal{I}$ is a tangent vector to q . An example is provided next.

3.4 Example: square root curvature fields

We introduce a new second-order feature map, which we call square root curvature field (SRCF). Combined with

the SRNF feature map, it yields a meaningful second-order elastic metric, whose geodesic distance can be computed efficiently similarly to (15). The SRCF is defined as

$$\Phi_q := H_q A_q^{1/2}, \quad (18)$$

where H_q is the vector-valued mean curvature of q . This feature map Φ simultaneously encodes information about the normal vector, the area form, and the curvature of the shape. The SRCF distance (16) coincides with the Willmore (1993) energy of q_1 if q_0 is a minimal surface. More generally, this distance quantifies differences in the mean curvatures of q_0 and q_1 , weighted symmetrically using the respective area forms of q_0 and q_1 . The associated SRCF metric (17) is a second-order elastic metric, whose highest-order term is the normal component of the Laplacian Δ_q on the surface q :

$$G_q(h, h) = \int_M |(\Delta_q h)^\perp + C(\nabla h)|^2 A_q,$$

where $C(\nabla h)$ denotes a first-order term, which is a contraction of ∇h with the metric, cometric, and second fundamental form; see Appendix D. The same form of metric is obtained using the scalar instead of the vector-valued mean curvature, but the feature map contains less information in this case. Despite the complicated form of the SRCF metric, its geodesic distance is easy to compute, thanks to the Hilbert structure on the range of the feature map. We also experimented with alternative forms of the SRCF, where the vector-valued mean curvature is replaced by the scalar mean curvature; see Section 4.

4 Implementation

This section describes the implementation of the geodesic boundary value problem, i.e., the surface matching algorithm for triangular meshes. The term geodesic shall be used quite liberally, even when approximations, relaxations, and discretizations of the Riemannian energy functional are involved.

4.1 Open source library

All algorithms described in this section are publicly available² as a Python package `SRNFmatch`. These include the routines for symmetric and asymmetric surface matching based on SRNF and SRCF energies, inverting the SRNF and SRCF maps, computing geodesic

² https://github.com/SRNFmatch/SRNFmatch_code

interpolations, and generating some of the figures in this paper. Our implementation relies on the following Python libraries: `Numpy`, `Scipy`, `Pytorch`, `PyKeops`, and `Vtk`. We refer to the online documentation for further details on the code.

4.2 Mesh discretization

To discretize the space of surfaces, we use a triangular surface as the domain M of the function space \mathcal{I} and restrict this function space to piece-wise linear functions. Alternative discretizations for more general mesh structures could be used, as well. Any such mesh structure defines a natural reference density on M , which assigns weight 1 to every face of the mesh. This reference density and its corresponding half-density are used throughout to identify densities and half-densities with functions. Alternative reference densities would lead to the same results but involve additional quadrature weights.

In the computation of SRNF or SRCF distances, the triangulation of M can be chosen freely but has to coincide for the initial and target surfaces. In contrast, in the geodesic boundary value problem, the initial and final surfaces may have different mesh structures and even different topologies. This flexibility is granted by the varifold terms.

4.3 SRNF distances between triangular meshes

The SRNF of a triangular mesh is the piece-wise constant function (1), which associates to any face of the mesh the square root of its area multiplied by the unit-normal vector. Thus, the SRNF of a surface $q \in \mathcal{I}$ can be written as a function $N_{q,f}$ of the faces f of the mesh:

$$N_{q,f} := n_{q,f} A_{q,f}^{1/2},$$

where $n_{q,f}$ and $A_{q,f}$ are the unit normal vector and area, respectively, of face f on surface q . This discrete definition of the SRNF is consistent with the continuous one: it preserves convergence almost everywhere of the surface q and its tangent map Tq , and it is uniquely determined by this property. Given two triangular meshes q_0 and q_1 with the same mesh structure, the corresponding SRNF distance is thus a sum over all faces:

$$\text{dist}_N(q_0, q_1)^2 = \sum_f |N_{q_0,f} - N_{q_1,f}|^2. \quad (19)$$

4.4 SRCF distances between triangular meshes

Higher-order feature maps such as the SRCF do not extend uniquely by a limiting procedure from smooth surfaces to triangular meshes. Therefore, we resort to principles of discrete differential geometry for their definition (see Bobenko et al. 2008). Recall that the SRCF of a smooth surface q is defined as $\Phi_q = H_q A_q^{1/2}$, where H_q is the vector mean curvature. Following Sullivan (2008), we define the discrete SRCV on each vertex v as

$$\Phi_{q,v} := \frac{\frac{1}{2} \sum_{e \ni v} e \times (n_1 - n_2)}{\sqrt{\frac{1}{3} \sum_{f \ni v} A_{q,f}}},$$

where e and f run through all edges and faces adjacent to the vertex v , and n_1, n_2 are the unit-normal vectors of the two faces adjacent to the edge e . Equivalently,

$$\Phi_{q,v} = \frac{\frac{1}{2} \sum_{f \ni v} e \times n}{\sqrt{\frac{1}{3} \sum_{f \ni v} A_{q,f}}}, \quad (20)$$

where e is the edge opposite of v in f , and n is the unit-normal vector of f . Both formulas can be expressed equivalently by the well-known co-tangent formula (Sullivan 2008). The discrete SRCF energy is defined by summing over all vertices:

$$\text{dist}_\Phi(q_0, q_1)^2 := \sum_v (\Phi_{q_0,v} - \Phi_{q_1,v})^2.$$

4.5 Varifold distances between triangular meshes

Varifold distances have a natural discretization for triangular meshes. We briefly recap the main elements using the notation of Section 2.3, referring to Kaltenmark et al. (2017) or Charon et al. (2020) for further details. For any triangular surfaces q_0 and q_1 , the varifold inner product (8) can be expressed as a sum over pairs of faces f_0 and f_1 :

$$\langle \mu_{q_0}, \mu_{q_1} \rangle_{V^*} = \sum_{f_0, f_1} \iint_{f_0 \times f_1} \rho(|x_0 - x_1|) \gamma(n_{q_0}(x_0) \cdot n_{q_1}(x_1)) A_{q_0}(x_0) A_{q_1}(x_1)$$

Importantly, meshes q_0 and q_1 may have different topologies. Note that the normal vector fields are constant on the domain of the double integral, i.e., on each pair of faces. Moreover, assuming that the faces are small compared to the variation scale of the kernel ρ , the function $\rho(|x_0 - x_1|)$ is well-approximated on the domain of the double integral by the constant $\rho(|c_{q_0, f_0} - c_{q_1, f_1}|)$, where $c_{q,f}$ is the barycenter of face f within surface q :

$$c_{q,f} := (\sum_{v \in f} v) / (\sum_{v \in f} 1)$$

The approximation we use then writes:

$$\langle \mu_{q_0}, \mu_{q_1} \rangle_{V^*} \approx \sum_{f_0, f_1} \rho(|c_{q_0, f_0} - c_{q_1, f_1}|) \gamma(n_{q_0, f_0} \cdot n_{q_1, f_1}) A_{q_0, f_0} A_{q_1, f_1}$$

This discretization easily extends to functional varifolds, as described in Section 2.3: the varifold inner product of textured triangular surfaces (q_0, ζ_{q_0}) and (q_1, ζ_{q_1}) is computed as

$$\langle \mu_{q_0, \zeta_{q_0}}, \mu_{q_1, \zeta_{q_1}} \rangle_{V^*} \approx \sum_{f_0, f_1} \rho(|c_{q_0, f_0} - c_{q_1, f_1}|) \gamma(n_{q_0, f_0} \cdot n_{q_1, f_1}) \tau(|\zeta_{q_0, f_0} - \zeta_{q_1, f_1}|) A_{q_0, f_0} A_{q_1, f_1},$$

where $\zeta_{q,f}$ denotes the texture value attached to face f of surface q . If the texture is given on vertices rather than faces, one first distributes it to faces by setting

$$\zeta_{q,f} := (\sum_{v \in f} \zeta_{q,v}) / (\sum_{v \in f} 1).$$

The accuracy of the above approximation can be precisely controlled in terms of the maximum diameter of faces and the modulus of continuity of the kernels ρ and γ , as shown by Kaltenmark et al. (2017). The approximation can be also interpreted as replacing the continuous varifold associated to the surface with a finite sum of Dirac masses. One can see that the numerical complexity of the varifold distance evaluation is quadratic in the number of faces.

4.6 Energy computation and optimization

For the minimization of the matching energies (9) and (10) over the vertex coordinates we use a quasi-Newton method, namely SciPy’s implementation of the L-BFGS algorithm introduced by Liu and Nocedal (1989). The most costly operation at each iteration of the optimization routine is the computation of the varifold distance and its gradient. Indeed, the computational complexity of varifold distances is quadratic in the number of faces, whereas the complexity of SRNF distances is only linear. Thus, efficient varifold computations are critical for the overall speed of our matching algorithm. Thanks to their highly parallelizable structure, they are well-suited for graphics processor units (GPUs). Our implementation leverages the Python library PyKeops, which has recently been developed by Charlier et al. (2020). This library is tailored specifically for the fast evaluation of kernel reductions using CUDA. Moreover, this library allows for automatic differentiation of all expressions. The SRNF distance is directly implemented in PyTorch and therefore also allows for automatic differentiation.

4.7 Multi-resolution

The boundary value problem naturally lends itself to a multi-resolution implementation, where point correspondences are initially computed on a coarse grid and then iteratively refined. This makes it easier to find an initial guess where the optimizer for the matching energy (9) or (10) does not get trapped in a local minimum. Similar multi-resolution approaches have been developed by Kilian et al. (2007) for given point correspondences and by Marin et al. (2019) in the context of functional maps. For simplicity, we only describe the implementation of the symmetric matching problem (10). This is detailed in Algorithm 1 and illustrated in Figure 2. The input to the algorithm are two triangular surfaces q_0 and q_1 , which may have different mesh connectivities and topologies. At initialization, the variables \tilde{q}_0 and \tilde{q}_1 are set to some low-resolution surfaces with the same mesh structure. For instance, one may use a coarse triangulation of a sphere or some other problem-specific shape prior. Alternatively, one may also use a down-sampled version of q_0 or q_1 , albeit at the cost of destroying the symmetry of the matching algorithm. In any case, the initialization has to be done such that \tilde{q}_0 and \tilde{q}_1 share the same mesh structure, as this is required in the computation of SRNF distances. Then, iteratively, the surfaces \tilde{q}_0 and \tilde{q}_1 are deformed in order to minimize the matching functional (10) and up-sampled to higher resolutions. The energy minimization is implemented as described in the previous section. Up-sampling is implemented using a subdivision scheme. The easiest choice is a non-adaptive scheme, which divides each triangle into four sub-triangles, but adaptive schemes would be possible, as well. Functional data, if present, have to be mapped to the new mesh based on the initial fully sampled surface by e.g. closest point(s) interpolation.

Algorithm 1: Multi-resolution matching

```

function Match( $\tilde{q}_0, \tilde{q}_1, q_0, q_1$ )
  Using an iterative optimization algorithm initialized
  with  $(\tilde{q}_0, \tilde{q}_1)$ , minimize the matching energy (10)
  with boundary data  $(q_0, q_1)$  over all surface meshes
  with the same combinatorics as  $(\tilde{q}_0, \tilde{q}_1)$  and return
  the minimizer.
end
function MultiResolutionMatch( $q_0, q_1$ )
  Initialize  $(\tilde{q}_0, \tilde{q}_1)$  at low resolution
   $(\tilde{q}_0, \tilde{q}_1) \leftarrow \text{Match}(\tilde{q}_0, \tilde{q}_1, q_0, q_1)$ 
  for several times do
     $(\tilde{q}_0, \tilde{q}_1) \leftarrow \text{UpSample}(\tilde{q}_0, \tilde{q}_1)$ 
     $(\tilde{q}_0, \tilde{q}_1) \leftarrow \text{Match}(\tilde{q}_0, \tilde{q}_1, q_0, q_1)$ 
  end
  return  $(\tilde{q}_0, \tilde{q}_1)$ 
end

```

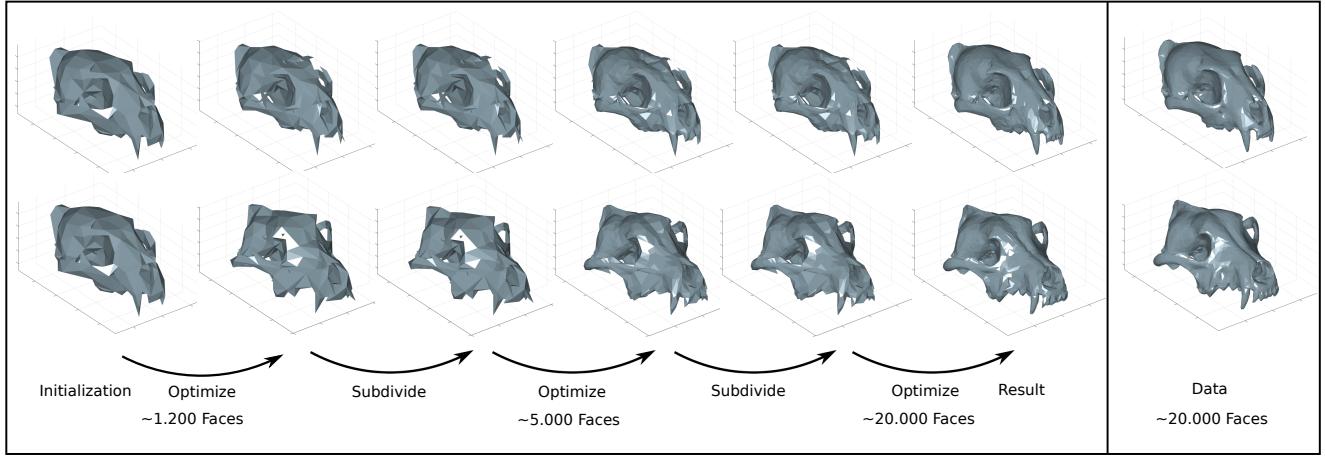


Fig. 2 Multi-resolution surface matching using Algorithm 1. The data q_0 and q_1 are high-resolution triangular meshes (last column). A down-sampled version of q_0 is used to initialize the surfaces \tilde{q}_0 and \tilde{q}_1 (first column). These surfaces are iteratively deformed to minimize the energy functional (10) and subdivided to achieve higher-resolution (middle columns). The result is a high-resolution point correspondence which solves the geodesic boundary value problem on shape space. The triangulated surfaces of the skulls are scans from the American Museum of Natural History in New York and have been obtained from Morphosource, a biological specimen database that has approximately 27,000 published 3D models of biological specimens.

4.8 Inversion of the SRNF map

Recall that Algorithm 1 computes the initial and final values of a geodesic between the given shapes. We next describe how to obtain intermediate values of this geodesic. These are of interest in many applications and describe the optimal elastic deformation between the shapes. Our algorithm is an extension and modification of Laga et al. (2017). Pseudo-code for the algorithm is shown in Algorithm 2. The starting point is the observation that geodesics in the feature space of the SRNF map are straight lines. Thus, given the output \tilde{q}_0 and \tilde{q}_1 of Algorithm 1, this geodesic in feature space can be computed by linearly interpolating the SRNFs of \tilde{q}_0 and \tilde{q}_1 . To find surfaces corresponding to these interpolated SRNFs, one has to invert the SRNF map. Of course, inversion in the strict sense is impossible because the SRNF map is neither injective nor surjective (Klassen and Michor 2020). In particular, the linear interpolation of two SRNFs may not be the SRNF of any surface, and if such a surface exists it may be non-unique. We circumvent this issue by recasting it as an optimization problem: given a function $\tilde{N}: M \rightarrow \mathbb{R}^3$, which is constant on each face f of the triangular surface M , find a triangular surface mesh q which minimizes the energy functional

$$\int_M |N_q - \tilde{N}|^2 = \sum_f |N_{q,f} - \tilde{N}_f|^2. \quad (21)$$

Similarly as in the matching problem, we minimize this energy functional using an L-BFGS method, where the gradient is computed by automatic differentiation in

PyTorch. This improves upon the optimization routine of Laga et al. (2017), where no information about the Hessian is used. Despite the analytical difficulties pointed out in Klassen and Michor (2020), this optimization problem turns out to be rather well-behaved in our experiments. For instance, the optimizer is able to recover a surface from its SRNF even when the initial guess is quite far off, as shown in Figure 3. Moreover, one actually disposes of quite accurate initial guesses in the context of matching problems, where a linear path in SRNF-space has to be inverted. An example of a resulting optimal elastic deformation is shown in Figure 4. A similar strategy can be used for the approximate inversion of the SRCF map, where one augments the energy functional (21) by the L^2 -difference between the SRCFs.

Algorithm 2: Computation of geodesics

```

function InverseSRNF( $\tilde{N}$ )
  Minimize the energy (21) over all triangular meshes
   $q$  with the same combinatorics as  $\tilde{N}$  and return the
  minimizer.
end
function GeodesicInterpolation( $t, q_0, q_1$ )
   $(\tilde{q}_0, \tilde{q}_1) \leftarrow \text{MultiResolutionMatch}(q_0, q_1)$ 
   $\tilde{N} \leftarrow (1-t)N_{\tilde{q}_0} + tN_{\tilde{q}_1}$ 
   $q \leftarrow \text{InverseSRNF}(\tilde{N})$ 
  return  $q$ 
end

```

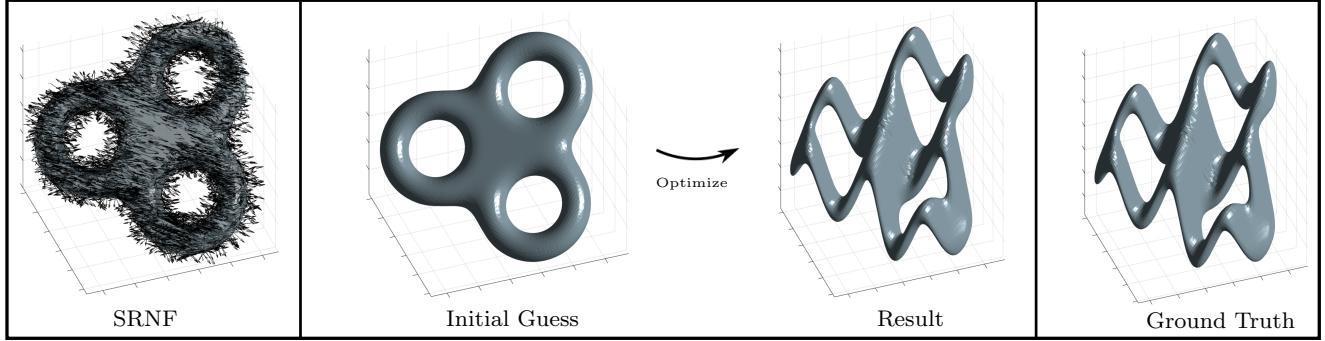


Fig. 3 Approximate inversion of the SRNF-map. Given the SRNF (left) of an unknown surface (right), optimization of the energy functional (21) determines a surface (middle) whose SRNF closely matches the given one. The algorithm works reliably for high-genus and high-resolution surfaces (here with approximatively 25,000 faces).

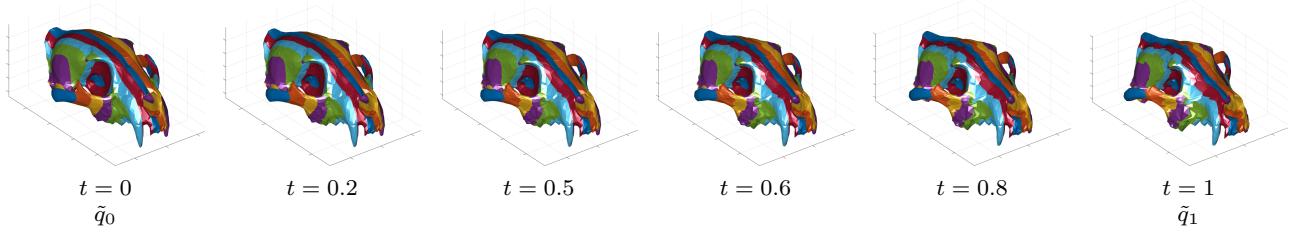


Fig. 4 Computation of geodesics using Algorithm 2. The geodesic is obtained by matching the given surfaces q_0 and q_1 using Algorithm 1, linearly interpolating the SRNFs of the matched surfaces \tilde{q}_0 and \tilde{q}_1 , and inverting the SRNF map. The color-coding shows that all surfaces along the geodesic are parametrized consistently.

5 Numerical results

In this final section we demonstrate various features of our algorithm through numerical simulations.

5.1 Choice of parameters

There are several parameters to be set. Specifically, these are the kernels in the varifold fidelity metric and the coefficient λ which weighs the contribution of the SRNF or SRCF term relative to the varifold terms in the matching energy. In all our simulations, we used the Gaussian kernel

$$\rho(|x_1 - x_2|) := \exp(-|x_2 - x_1|^2/\sigma^2)$$

of scale σ on \mathbb{R}^3 and the Binet kernel

$$\gamma(n_1 \cdot n_2) := (n_1 \cdot n_2)^2$$

on S^2 . Our implementation allows for many other choices of kernels, and we refer to Kaltenmark et al. (2017) for a thorough discussion of the effects that this can have in shape matching problems. The kernel scale σ and weight λ have to be selected in accordance with the data. This can be done in an adaptive fashion in combination with the multi-resolution approach described earlier. Namely, we typically initialize the kernel scale

σ to a value around a tenth of the shape's diameter and decrease it throughout the successive runs of the multi-resolution algorithm. On the opposite, the relaxation parameter λ is increased after each run in order to enforce a closer and closer matching to the target shape. We also point to the demo scripts provided in the [github](#) package for examples of parameter setting.

5.2 Choice of elastic energy

We used either the SRNF distance or a combination of the SRNF and SRCF distances as deformation energies in our numerical experiments. SRNF distances reliably led to good results and stable optimization routines and were used for most of our numerical simulations. SRCF distances alone had the tendency to create numerical instabilities. Combined with SRNF distances, according to our limited numerical evidence, these instabilities disappeared.

As expected, the choice of deformation energy influences the resulting point correspondences and distances, and thus impacts the statistical analysis on the space of shapes. This is demonstrated in a toy example in Figure 6. In future works it will be interesting to choose the deformation energy in a data-driven way, as suggested by Needham and Kurtek (2020) for elastic

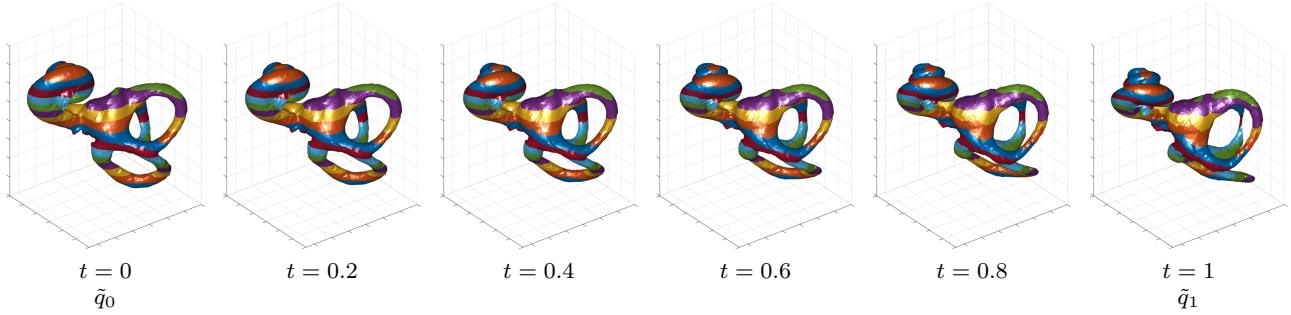


Fig. 5 Matching high-genus surfaces. Two cochleas (left, right) are given as triangular meshes without any point correspondence between them. Algorithms 1 and 2 determine a consistent parametrization (color-coded) for the two surfaces and the geodesic in-between (middle).

metrics on spaces of curves and by Niethammer et al. (2019) for diffeomorphic metrics.

5.3 Run time

As an indication of the algorithm’s practical performance, we show in Table 1 the average run time per iteration of the L-BFGS optimization scheme for the (non-symmetric) surface matching algorithm. Typically, in those experiments, the algorithm required around 300 iterations to reach a relative error of less than 0.1%. All those computation times were obtained on an Intel i7-9700K processor with 3.6 GHz clock rate and a 2018 Nvidia GeForce RTX 2080 Ti with 11GB memory. We ran the matching between two surfaces with a range of different mesh sizes, both for the SRNF and the combined SRNF-SRCF matching procedures.

To further put those results into perspective, we also performed similar tests for another well-established surface registration framework based on the Large Deformation Diffeomorphic Metric Mapping (LDDMM) model of Beg et al. (2005). This approach shares some similarity with the one presented in this paper in that both solve inexact matching problems with varifold distances as relaxation terms, the essential difference being that the LDDMM algorithm estimates a dense and smooth deformation of the whole ambient space to transform the source shape. This typically induces a higher numerical complexity compared to the SRNF framework that directly optimizes over the deformed source shape, as evidenced by run times per iteration approximately 10 times larger. For the computation of the diffeomorphic flow, we used a Ralston integrator with 10 time steps. We emphasize that, in order to make those comparisons meaningful, both the SRNF and LDDMM models were implemented in Python based on the Scipy implementation of the L-BFGS algorithm and rely on CUDA subroutines as well as automatic differentiation

through the use of the same PyTorch and PyKeops libraries.

Faces	Time per iteration (s)		
	SRNF	SRNF & SRCF	LDDMM
1,000	0.007	0.01	0.11
10,000	0.016	0.018	0.22
50,000	0.12	0.12	1.1
200,000	1.73	1.84	16.1

Table 1 Computation times for a single iteration of the L-BFGS algorithm, which is here applied to minimize the non-symmetric matching energy (9), for different resolutions of the mesh. For comparison, we also provide the similar time per iteration for an alternative surface registration approach based on a diffeomorphic model (LDDMM).

5.4 High-genus surfaces

All parts of the algorithm, from varifold distances to SRNF distances and SRNF inversion, are implemented for triangular meshes. Therefore, there are no topological restrictions on the surfaces to be matched. For instance, it is possible to match high-genus surfaces such as skulls or cochleas, as demonstrated in Figures 4 and 5. In these examples, the surfaces q_0 and q_1 to be matched are given by triangular meshes with different combinatorics, and no point correspondences are known a priori. The matching algorithm is initialized with a down-sampled version of q_0 , as described in Section 4.7. Accordingly, the resulting interpolated surfaces have the same genus as q_0 . The deformation energy, which was used in these examples, is the SRNF energy. The multi-resolution matching algorithm converged in less than 20 seconds for both examples. The resulting point correspondences are color-coded in the figures.

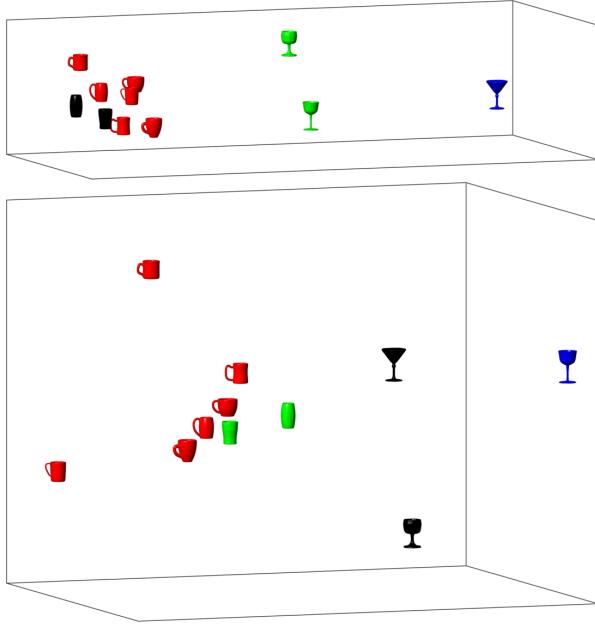


Fig. 6 Matching surfaces with different topologies. The cups are positioned according to multi-dimensional scaling based on the matrix of pairwise elastic-varifold discrepancies, and are colored according to an automatic distance-based clustering algorithm. Positions and colors consistently indicate that the elastic-varifold distance is able to quantify and detect both topological and geometric variations. The distances were computed using only SRNFs (top) or SRNFs together with SRCFs (bottom).

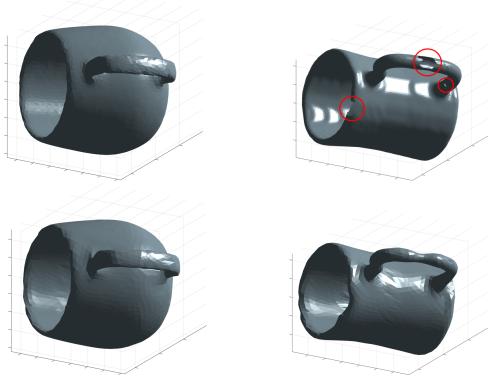


Fig. 7 Matching surfaces with topological noise. One of the two cups to be matched (top) has several artifacts, including a broken-off piece at the top and some holes in the triangulation at the handle. The matching algorithm is initialized by a surface without these artifacts and produces a decent fit (bottom).

5.5 Topological inconsistencies

Our algorithm can handle surfaces with different topologies. This would not be possible in a pure elastic-matching approach, where the surfaces are required to be homotopic, see e.g. Jermyn et al. (2017). This is possible in our approach thanks to the combination with

varifold distances. The matching energy (10), which is a sum of elastic and varifold terms, quantifies both geometric and topological differences and can therefore be used as a measure of discrepancy between shapes of arbitrary topology. For example, as shown in Figure 6, it is possible to compare cups with handles to glasses without handles. In this comparison, the presence or absence of handles as well as differences in the shape of the glasses or cups are all taken into account, as revealed by clustering based on the matrix of pairwise elastic-varifold discrepancies.

The initialization of the matching algorithm involves an important choice to be made, namely, fixing a topology and mesh structure for the surfaces \tilde{q}_0 and \tilde{q}_1 (see Algorithm 1). In our experiments we initialized these surfaces once with the topology of the source and once with the topology of the target and used the average of the two resulting distances as our final discrepancy measure.

5.6 Topological noise

Handling topological noise is a further application of the algorithm's capability of dealing with different topologies. An example is presented in Figure 7, which shows enlarged pictures of two of the cups in Figure 6. One of these cups has several topological artifacts: a solid piece at the top of the cup has broken off, and several triangles are missing in the mesh describing the cup's handle. Nevertheless, the cup can be matched to a similar cup without these artifacts. Of course, for topological reasons, the match cannot be exact, but it nevertheless provides a very good approximation without the missing triangles.

Topological artifacts are often created by 3D scanning technology, as e.g. in the case of the facial scans in Figure 8. In this case, both facial surfaces to be matched suffer from topological noise. The first surface has more missing areas than the second one and is used to initialize the matching algorithm. The resulting geodesic interpolation gradually deforms the first face into the second one while filling out the missing areas.

5.7 Functional shapes

Our algorithm provides a convenient way of handling texture information, as explained in Section 2.3. Namely, the varifold distance can be augmented by an additional term which measures texture differences. Texture information can guide the optimizer in the matching problem and improve the resulting point correspondences.

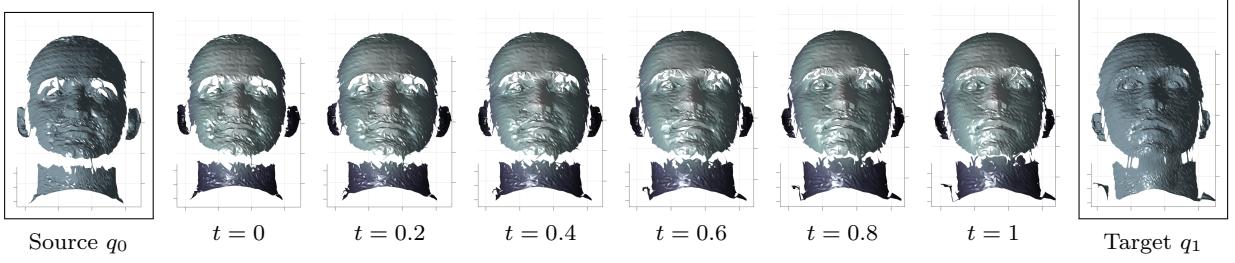


Fig. 8 Matching surfaces with topological noise. The facial scanning technology produced meshes (left, right) of varying degrees of incompleteness. Nevertheless, these meshes are matched correctly, as shown by the geodesic interpolation between them (middle).

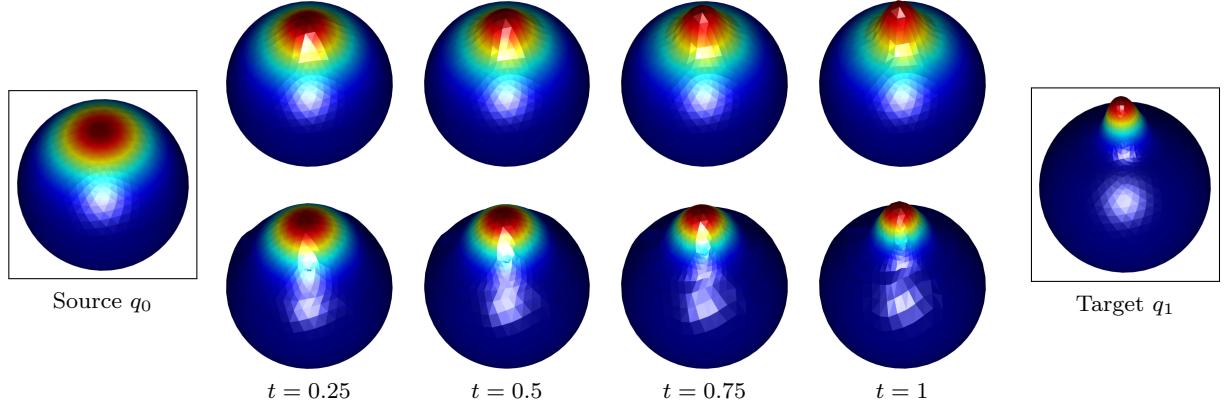


Fig. 9 Influence of texture on optimal point correspondences. Matching the textures requires compression of a large portion of a sphere q_0 (left) to a small portion of a deformed sphere q_1 (right). This compression occurs in the geodesic between the textured shapes (bottom) but not between the untextured shapes (top). It increases the SRNF distance between the matched surfaces \tilde{q}_0 and \tilde{q}_1 from 0.76 to 1.74, but this increase is compensated by a better fit of textures.

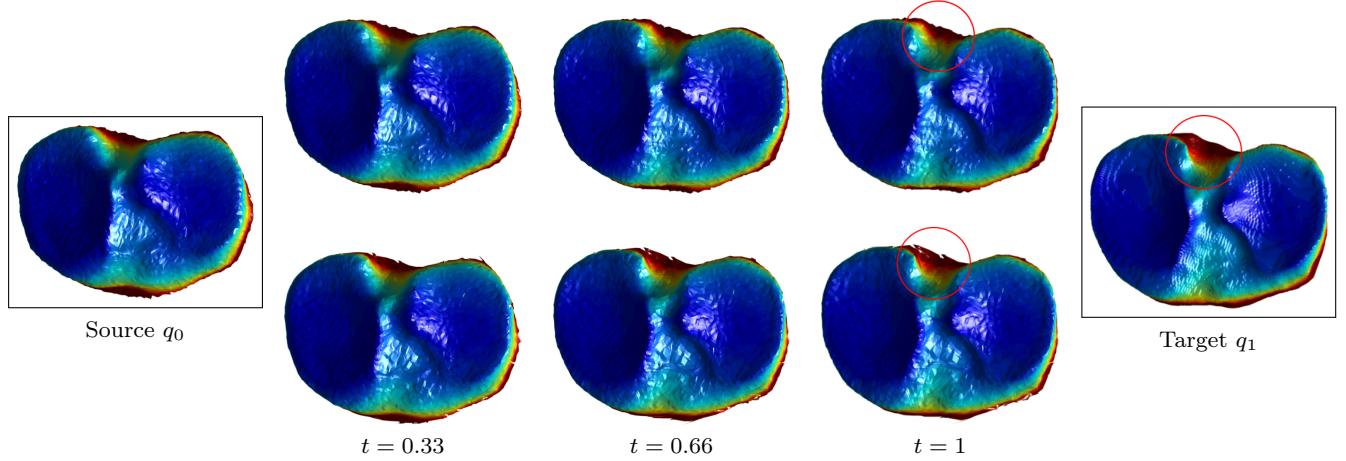


Fig. 10 Influence of texture on optimal point correspondences. Two tibia bones (left, right) carry texture information representing the width of the tibia-femoral joint. Differences in width are most pronounced in the upper part of the bone. This area is matched better by the geodesic interpolation of textured shapes (bottom) than of untextured shapes (top), see particular the area highlighted by the red circles.

We demonstrate this in two examples, first on synthetic and then on real-world data.

As a synthetic example, we consider in Figure 9 texture information which asks for compression of a large portion of a sphere to a small portion of a deformed sphere. The texture signals range from 0 to 1, and the texture kernel τ is Gaussian with a scale of 0.2. Without any texture information, there is no compression in the geodesic between the two shapes. If, however, texture information is included in the matching energy, then there is compression along the geodesic, in accordance with the texture information. Visually, the difference is most prominent in the terminal values of the two geodesics. Quantitatively, the difference can also be seen from the resulting SRNF distance, which is higher when texture information is included (1.74 versus 0.76). The reason is that the optimizer has to strike a balance between matching the textures and minimizing the deformation energy.

The same qualitative behavior can be seen in the real-world example in Figure 10. Here the data consist of tibia bone extremities of two subjects, and the texture represents the width of the tibia-femoral joint, estimated using the approach of Cao et al. (2015). The texture signals range from 0 to 128, and the texture kernel is Gaussian with three successive scales of 100, 50, and 25. Differences in texture are most prominent in the upper middle part of the bone, where the red region is less pronounced in the left than in the right sample. These regions are matched incorrectly when no texture information is used (top), but are matched more accurately otherwise (bottom).

5.8 Population statistics

Finally we show how our surface matching algorithm can be used to perform simple statistical analyses of shape populations, following the framework of geometric statistics; see for instance the recent review in Pennec et al. (2019). A first example is multi-dimensional scaling or k-means clustering based on the pairwise distance matrix; see Figure 6. As a second example, we consider the estimation of the Karcher mean of a dataset of surface meshes. Figure 11 shows the Karcher mean of 18 amygdala shapes from the BIOCARD database (Miller et al. 2015) which were segmented from MRI images. To approximate the Karcher mean we followed a method adapted from Arnaudon and Nielsen (2013) and Cury et al. (2013), which consists in an iterative geodesic centroid procedure. The details of the algorithm are described in Algorithm 3.

Algorithm 3: Karcher mean

```

function Karcher( $q_1, \dots, q_n$ )
  Initialize  $\bar{q}$  to some mesh, which determines the
  topology/connectivity of the mean
  for  $i \leftarrow 1$  to some number of iterations do
     $(q_1, \dots, q_n) \leftarrow \text{RandomShuffle}(q_1, \dots, q_n)$ 
    for  $j \leftarrow 1$  to  $n$  do
       $(\tilde{q}, \tilde{q}_j) \leftarrow \text{MultiResolutionMatch}(q, q_j)$ 
       $\bar{N} \leftarrow N_{\tilde{q}} + \frac{1}{(i-1)*n+j} (N_{\tilde{q}_j} - N_{\tilde{q}})$ 
       $\bar{q} \leftarrow \text{InverseSRNF}(\bar{N})$ 
    end
  end
  return  $\bar{q}$ 
end

```

6 Discussion and conclusion

The above examples demonstrate that our implementation is fast and versatile, and that it produces quite natural distances and interpolations between surfaces, including surfaces which have different topologies or carry texture information. These advantages are the result of our novel combination of elastic metrics and varifold distances. We also introduced the possibility of extending this idea to other surface feature maps, in particular square root curvature fields, as a way to emulate higher-order Sobolev metrics on surfaces. However, a precise study of the properties and effects of square root curvature fields in surface matching is left for future work.

Besides these advantages, we want to mention some important caveats of our SRNF-based approach to surface matching. On the one hand, as the SRNF metric is only of first order and incomplete, minimizing the SRNF energy (5) over reparametrizations can lead to shrinkage of some parts of a surface. Accordingly, “optimal reparametrizations” may in fact be singular. On the other hand, varifold distances are by construction insensitive to structures of vanishing or small area, such as spikes or one-dimensional fibers. Thus, the combination of these two distances in our proposed formulation makes the optimizer in our matching procedure prone to finding surfaces with spiking singularities, despite the SRNF distance still being computed accurately. These degeneracies are alleviated considerably by the multi-resolution minimization scheme which we introduced. Nevertheless, they can still be routinely observed in some simulations, most notably around irregular boundaries of objects.

An important question for future work is whether this problem can be solved by suitable regularization, either using higher-order elastic metrics (e.g. based on SRCFs) or higher-order fidelity metrics (e.g. normal cycles), which may not have the same vanishing behavior as varifolds. More generally, higher-order feature maps

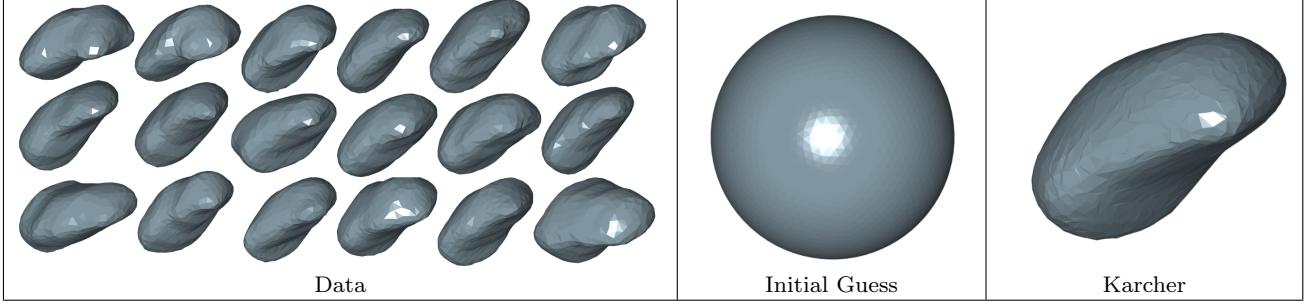


Fig. 11 The Karcher mean (right) of a dataset of amygdala shapes (left), computed using Algorithm 3 with a sphere (center) as initialization.

might also be of theoretical interest in order to overcome the incompleteness of low-order metrics.

Acknowledgements The authors thank Stanley Durrleman, José Braga, and Jean Dumoncel for the use of the cochlea data, Wojtek Zbijewski and his group for sharing the tibia surfaces as well as Daniel Tward and the BIOCARD team for sharing the amygdala dataset. In addition, we thank Zhe Su for his help with Figure 1 and the whole shape group at Florida State University for helpful discussions during the preparation of this manuscript.

A Notation

Throughout this section, $h \in T_q \mathcal{I} = C^\infty(M, \mathbb{R}^3)$ denotes a tangent vector to $q \in \mathcal{I}$, and X, Y are vector fields on M . Traces are denoted by Tr or a dot, T is the tangent functor, and D stands for directional derivatives. For instance, the derivative at q in the direction of h is denoted by $D_{(q,h)}$. We write $\bar{g} = \langle \cdot, \cdot \rangle$ for the Euclidean inner product on \mathbb{R}^3 , $|\cdot|$ for the Euclidean norm on \mathbb{R}^3 , $g_q = q^* \bar{g}$ for the pull-back metric on TM , g_q^{-1} for the cometric on T^*M , and $g_q^{-1} \otimes \bar{g}$ for the product metric on $T^*M \otimes \mathbb{R}^3$. The metric g_q corresponds to a fiber-linear map \flat from TM to T^*M , and the cometric g_q^{-1} corresponds to a fiber-linear map \sharp from T^*M to M . The Riemannian surface measure of g_q is denoted by A_q , and the corresponding half-density by $A_q^{1/2}$. The surface measure is a section of the volume bundle Vol , and the half-density of the half-density bundle $\text{Vol}^{1/2}$. The normal projection $\perp: M \rightarrow L(\mathbb{R}^3, \mathbb{R}^3)$ is defined as $\perp = \langle \cdot, n_q \rangle n_q$, the tangential projection $\top: M \rightarrow L(\mathbb{R}^3, TM)$ is defined as $\top = (Tq)^{-1}(\text{Id}_{\mathbb{R}^3} - \perp)$, and one has the identity $\perp + Tq \circ \top = \text{Id}_{\mathbb{R}^3}$. Depending on the context, ∇ is the covariant derivative on \mathbb{R}^3 , which coincides with the usual coordinate derivative, or the Levi-Civita covariant derivative of g_q . For instance, in the definition $\nabla_{X,Y}^2 h := \nabla_X \nabla_Y h - \nabla_{\nabla_X Y} h$, only $\nabla_X Y$ is the Levi-Civita covariant derivative, and all other derivatives are coordinate derivatives.

B Formula for SRNF metrics

In this section we establish the explicit formula (14) for the SRNF metric. We need some variational formulas from e.g.

Bauer et al. (2012):

$$\begin{aligned} D_{(q,h)} A_q &= \text{Tr}(g_q^{-1} \cdot \langle \nabla h, Tq \rangle) A_q = \text{Tr}((\nabla h)^\top), \\ D_{(q,h)} A_q^{1/2} &= \frac{1}{2} \text{Tr}((\nabla h)^\top) A_q^{1/2}, \\ D_{(q,h)} n_q &= -Tq \circ \langle n_q, \nabla h \rangle^\sharp, \\ D_{(q,h)} N_q &= \left(-Tq \circ \langle n_q, \nabla h \rangle^\sharp + n_q \frac{1}{2} \text{Tr}((\nabla h)^\top) \right) A_q^{1/2}. \end{aligned}$$

Putting these together, one obtains the following expression for the SRNF metric (13):

$$\begin{aligned} G_q(h, h) &:= \int_M |D_{(q,h)} N_q|^2 \\ &= \int_M |Tq \circ \langle n_q, \nabla h \rangle^\sharp|^2 A_q \\ &\quad + \int_M |n_q \frac{1}{2} \text{Tr}((\nabla h)^\top)|^2 A_q \\ &= \int_M |\langle n_q, \nabla h \rangle|_{g_q^{-1}}^2 A_q + \frac{1}{4} \int_M \text{Tr}((\nabla h)^\top)^2 A_q \\ &= \int_M |(\nabla h)^\perp|_{g_q^{-1} \otimes \bar{g}}^2 A_q + \frac{1}{4} \int_M \text{Tr}((\nabla h)^\top)^2 A_q. \end{aligned}$$

C Approximation properties of SRNF distances

This section makes precise in what sense the SRNF distance approximates the geodesic distance of the SRNF metric. The geodesic distance of the SRNF metric is lower-bounded by the SRNF distance because the latter is a chordal distance:

$$\begin{aligned} \text{dist}_{\mathcal{I}}(q_0, q_1)^2 &:= \inf_q \int_0^1 G_q(\partial_t q, \partial_t q) dt \\ &= \inf_q \int_0^1 |D_{(q, \partial_t q)} N_q|^2 dt \geq \|N_{q_1} - N_{q_0}\|_{L^2}^2, \end{aligned}$$

where the infimum is over all paths q as in (11). Conversely, the geodesic distance of the SRNF metric is upper-bounded by the length of the linear interpolation between the immersions q_0 and q_1 , provided they are sufficiently close to each other for this to make sense, leading to the upper bound

$$\begin{aligned} \text{dist}_{\mathcal{I}}(q_0, q_1)^2 &\leq G_{q_0}(q_1 - q_0, q_1 - q_0) \\ &= \|D_{(q_0, q_1 - q_0)} N_{q_0}\|_{L^2}^2 \approx \|N_{q_1} - N_{q_0}\|_{L^2}^2, \end{aligned}$$

which is valid in any chart around q_0 up to terms of order $o(\|N_{q_1} - N_{q_0}\|_{L^2}^2)$.

D Formula for SRCF metrics

Recall that the scalar and vector-valued second fundamental forms are defined as

$$s_q(X, Y) := \langle \nabla_X(Tq \circ Y), n_q \rangle, \quad S_q(X, Y) := s_q(X, Y)n_q,$$

for any vector fields X and Y on M . Following Bauer et al. (2012), one obtains the variational formulas

$$\begin{aligned} D_{(q,h)}s_q(X, Y) &= \langle D_{(q,h)}\nabla_X(Tq \circ Y), n_q \rangle \\ &\quad + \langle \nabla_X(Tq \circ Y), D_{(q,h)}n_q \rangle \\ &= \langle \nabla_X\nabla_Y h, n_q \rangle - \langle \nabla_X(Tq \circ Y), Tq \circ \langle \nabla h, n_q \rangle^\sharp \rangle \\ &= \langle \nabla_X\nabla_Y h, n_q \rangle - g_q(\nabla_X Y, \langle \nabla h, n_q \rangle^\sharp) \\ &= \langle \nabla_X\nabla_Y h, n_q \rangle - \langle \nabla_{\nabla_X Y} h, n_q \rangle = \langle \nabla_{X,Y}^2 h, n_q \rangle, \\ D_{(q,h)}S_q(X, Y) &= (D_{(q,h)}s_q(X, Y))n_q + s_q(X, Y)(D_{(q,h)}n_q) \\ &= (\nabla_{X,Y}^2 h)^\perp - s_q(X, Y)Tq \circ \langle \nabla h, n_q \rangle^\sharp. \end{aligned}$$

Using the formula $\Delta_q = -\text{Tr}(g_q^{-1}\nabla^2)$ for the Laplacian, this yields the following variational formula for the vector-valued mean curvature $H_q = \text{Tr}(g_q^{-1}S_q)$:

$$\begin{aligned} D_{(q,h)}g_q &= D_{(q,h)}\langle Tq, Tq \rangle = \langle \nabla h, Tq \rangle + \langle Tq, \nabla h \rangle, \\ D_{(q,h)}g_q^{-1} &= -g_q^{-1}(D_{(q,h)}g_q)g_q^{-1}, \\ D_{(q,h)}H_q &= \text{Tr}(g_q^{-1}D_{(q,h)}S_q) + \text{Tr}(S_qD_{(q,h)}g_q^{-1}) \\ &= -(\Delta_q h)^\perp - \text{Tr}(g_q^{-1}s_q)Tq \circ \langle \nabla h, n_q \rangle^\sharp \\ &\quad - 2\text{Tr}(S_qg_q^{-1}\langle \nabla h, Tq \rangle g_q^{-1}). \end{aligned}$$

Letting $C(\nabla h)$ denote the first-order terms in $-D_{(q,h)}H_q$, one obtains the desired formula for the SRCF metric:

$$G_q(h, h) := \int_M |D_{(q,h)}H_q|^2 A_q = \int_M |(\Delta_q h)^\perp + C(\nabla h)|^2 A_q.$$

References

Abe K, Erbacher J (1975) Isometric immersions with the same gauss map. *Mathematische Annalen* 215(3):197–201

Almgren F (1966) Plateau's Problem: An Invitation to Vari-fold Geometry. Student Mathematical Library

Arnaudon M, Nielsen F (2013) On approximating the Riemannian 1-center. *Computational Geometry* 46(1):93–104

Bauer M, Harms P, Michor PW (2011) Sobolev metrics on shape space of surfaces. *J Geom Mech* 3(4):389–438

Bauer M, Harms P, Michor PW (2012) Almost local metrics on shape space of hypersurfaces in n-space. *SIAM J Imaging Sci* 5(1):244–310

Bauer M, Bruveris M, Michor PW (2014) Overview of the geometries of shape spaces and diffeomorphism groups. *J Math Imaging Vision* 50(1-2):60–97

Bauer M, Bruveris M, Michor PW (2016) Why use Sobolev metrics on the space of curves. In: *Riemannian Computing in Computer Vision*, Springer, pp 233–255

Bauer M, Bruveris M, Charon N, Møller-Andersen J (2017) Varifold-based matching of curves via sobolev-type riemannian metrics. In: *Graphs in Biomedical Image Analysis, Computational Anatomy and Imaging Genetics*, Springer, pp 152–163

Bauer M, Bruveris M, Charon N, Møller-Andersen J (2019a) A relaxed approach for curve matching with elastic metrics. *ESAIM: Control, Optimisation and Calculus of Variations* 25:72

Bauer M, Charon N, Harms P (2019b) Inexact elastic shape matching in the square root normal field framework. In: *International Conference on Geometric Science of Information*, Springer, pp 13–20

Bauer M, Harms P, Michor PW (2020) Fractional Sobolev metrics on spaces of immersions. *Calculus of Variations and Partial Differential Equations* 59(62)

Beg MF, Miller MI, Trouvé A, Younes L (2005) Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* 61(2):139–157

Bernal J, Dogan G, Hagwood CR (2016) Fast dynamic programming for elastic registration of curves. In: *Computer Vision and Pattern Recognition (CVPR)*, pp 111–118

Bhattacharya A, Bhattacharya R (2012) Nonparametric inference on manifolds: with applications to shape spaces, vol 2. Cambridge University Press

Bobenko A, Schröder P, Sullivan JM, Ziegler GM (eds) (2008) *Discrete differential geometry*. No. 38 in Oberwolfach Semin., Birkhäuser

Bruveris M (2015) Completeness properties of Sobolev metrics on the space of curves. *Journal of Geometric Mechanics* 7(2):125–150

Bruveris M, Michor PW, Mumford D (2014) Geodesic completeness for Sobolev metrics on the space of immersed plane curves. In: *Forum of Mathematics, Sigma*, Cambridge University Press, vol 2

Cao Q, Thawait G, Gang GJ, Zbijewski W, Reigel T, Brown T, Corner B, Demehri S, Siewersden JH (2015) Characterization of 3D joint space morphology using an electrostatic model (with application to osteoarthritis). *Physics in medicine and biology* 60(3):947–60

Cervera V, Mascaró F, Michor PW (1991) The action of the diffeomorphism group on the space of immersions. *Differential Geom Appl* 1(4):391–401

Charlier B, Charon N, Trouvé A (2017) The fshape framework for the variability analysis of functional shapes. *Found Comput Math* 17(2):287–357

Charlier B, Feydy J, Glaunès J, Collin FD, Durif G (2020) Kernel operations on the GPU, with autodiff, without memory overflows. *arXiv preprint arXiv:200411127*

Charon N, Trouvé A (2013) The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM J Imaging Sci* 6(4):2547–2580

Charon N, Trouvé A (2014) Functional currents: a new mathematical tool to model and analyse functional shapes. *J Math Imaging Vision* 48(3):413–431

Charon N, Charlier B, Glaunès J, Gori P, Roussillon P (2020) Fidelity metrics between curves and surfaces: currents, varifolds, and normal cycles. In: *Riemannian Geometric Statistics in Medical Image Analysis*, Academic Press, pp 441 – 477

Cury C, Glaunes JA, Colliot O (2013) Template estimation for large database: a diffeomorphic iterative centroid method using currents. In: *International Conference on Geometric Science of Information*, Springer, pp 103–111

Dryden IL, Mardia KV (1998) Statistical shape analysis. Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons Ltd., Chichester

Ebin DG (1970) The manifold of Riemannian metrics. In: *Global Analysis (Proc. Sympos. Pure Math., Vol. XV, Berkeley, Calif., 1968)*, Amer. Math. Soc., Providence, R.I., pp 11–40

Federer H (1969) *Geometric measure theory*. Springer

Floater MS, Hormann K (2005) Surface parameterization: a tutorial and survey. In: *Advances in multiresolution for*

geometric modelling, Springer, pp 157–186

Frenkel M, Basri R (2003) Curve matching using the fast marching method. In: International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition, Springer, pp 35–51

Fröhlich S, Botsch M (2011) Example-driven deformations based on discrete shells. In: Computer graphics forum, Wiley Online Library, vol 30, pp 2246–2257

Geirhos R, Rubisch P, Michaelis C, Bethge M, Wichmann FA, Brendel W (2018) ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. arXiv preprint arXiv:181112231

Glaunès J, Qiu A, Miller MI, Younes L (2008) Large deformation diffeomorphic metric curve mapping. *Int J Comput Vis* 80(3):317

Grzegorzek M, Theobalt C, Koch R, Kolb A (2013) Time-of-Flight and Depth Imaging. Sensors, Algorithms and Applications: Dagstuhl Seminar 2012 and GCPR Workshop on Imaging New Modalities, vol 8200. Springer

Jermyn IH, Kurtek S, Klassen E, Srivastava A (2012) Elastic shape matching of parameterized surfaces using square root normal fields. In: European Conference on Computer Vision, Springer, pp 804–817

Jermyn IH, Kurtek S, Laga H, Srivastava A (2017) Elastic shape analysis of three-dimensional objects. *Synthesis Lectures on Computer Vision* 12(1):1–185

Kaltenmark I, Charlier B, Charon N (2017) A general framework for curve and surface comparison and registration with oriented varifolds. In: Computer Vision and Pattern Recognition (CVPR)

Kendall DG, Barden D, Carne TK, Le H (1999) Shape and shape theory. Wiley Series in Probability and Statistics, John Wiley & Sons Ltd., Chichester, DOI 10.1002/9780470317006

Kilian M, Mitra NJ, Pottmann H (2007) Geometric modeling in shape space. *ACM Trans Graphics* 26(3)

Klassen E, Michor PW (2020) Closed surfaces with different shapes that are indistinguishable by the srnf. *Archivum Mathematicum* 56(2):107–114

Kokoszka P, Reimherr M (2017) Introduction to functional data analysis. CRC Press

Kurtek S, Klassen E, Gore JC, Ding Z, Srivastava A (2012) Elastic geodesic paths in shape space of parameterized surfaces. *IEEE Trans Pattern Anal Mach Intell* 34(9):1717–1730

Laga H, Xie Q, Jermyn IH, Srivastava A (2017) Numerical inversion of SRNF maps for elastic shape analysis of genus-zero surfaces. *IEEE Trans Pattern Anal Mach Intell* 39(12):2451–2464

Lahiri S, Robinson D, Klassen E (2015) Precise matching of PL curves in \mathbb{R}^N in the square root velocity framework. *Geometry, Imaging and Computing* 2(3):133–186

Liu DC, Nocedal J (1989) On the limited memory bfgs method for large scale optimization. *Mathematical programming* 45(1-3):503–528

Marin R, Melzi S, Rodolà E, Castellani U (2019) High-resolution augmentation for automatic template-based matching of human models. In: 2019 International Conference on 3D Vision (3DV), IEEE, pp 230–239

Mennucci A, Yezzi A, Sundaramoorthi G (2008) Properties of Sobolev-type metrics in the space of curves. *Interfaces Free Bound* 10(4):423–445

Michor PW (2008) Topics in differential geometry, Graduate Studies in Mathematics, vol 93. American Mathematical Society, Providence, RI

Michor PW, Mumford D (2005) Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms. *Doc Math* 10:217–245

Michor PW, Mumford D (2006) Riemannian geometries on spaces of plane curves. *Journal of the European Mathematical Society* 8:1–48

Michor PW, Mumford D (2007) An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Appl Comput Harmon Anal* 23(1):74–113

Miller M, Ratnanather JT, Tward DJ, Brown T, Lee D, Ketcha M, Mori K, Wang MC, Mori S, Albert M, Younes L, BRT (2015) Network Neurodegeneration in Alzheimer’s Disease via MRI Based Shape Diffeomorphometry and High-Field Atlasing. *Frontiers in Bioengineering and Biotechnology* 3:54

Minh HQ, Murino V, Minh HQ (2016) Algorithmic Advances in Riemannian Geometry and Applications. Springer

Needham T, Kurtek S (2020) Simplifying Transforms for General Elastic Metrics on the Space of Plane Curves. *SIAM Journal on Imaging Sciences* 13(1):445–473

Niethammer M, Kwitt R, Vialard FX (2019) Metric learning for image registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp 8463–8472

Pennec X, Sommer S, Fletcher T (2019) Riemannian Geometric Statistics in Medical Image Analysis. Academic Press

Roussillon P, Glaunes JA (2016) Kernel metrics on normal cycles and application to curve matching. *SIAM J Imaging Sci* 9(4):1991–2038

Rumpf M, Wardetzky M (2014) Geometry processing from an elastic perspective. *GAMM-Mitteilungen* 37(2):184–216

Rumpf M, Wirth B (2015a) Variational methods in shape analysis. *Handbook of Mathematical Methods in Imaging* 2:1819–1858

Rumpf M, Wirth B (2015b) Variational time discretization of geodesic calculus. *IMA Journal of Numerical Analysis* 35(3):1011–1046

Sebastian TB, Klein PN, Kimia BB (2003) On aligning curves. *IEEE transactions on pattern analysis and machine intelligence* 25(1):116–125

Sheffer A, Praun E, Rose K, et al. (2007) Mesh parameterization methods and their applications. *Foundations and Trends in Computer Graphics and Vision* 2(2):105–171

Srivastava A, Klassen EP (2016) Functional and shape data analysis. Springer

Srivastava A, Klassen E, Joshi SH, Jermyn IH (2011) Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans Pattern Anal Mach Intell* 33(7):1415–1428

Su Z, Bauer M, Gallivan KA, Klassen E (2020a) Simplifying transformations for a family of elastic metrics on the space of surfaces. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)

Su Z, Bauer M, Preston SC, Laga H, Klassen E (2020b) Shape analysis of surfaces using general elastic metrics. *Journal of Mathematical Imaging and Vision* 62:1087–1106

Sullivan JM (2008) Curvatures of smooth and discrete surfaces. In: Bobenko et al. (2008), pp 175–188

Sundaramoorthi G, Mennucci A, Soatto S, Yezzi A (2011) A new geometric metric in the space of curves, and applications to tracking deforming objects by prediction and filtering. *SIAM J Imaging Sci* 4(1):109–145

Tumpach AB (2016) Gauge invariance of degenerate riemannian metrics. *Notices of the AMS* 63(4)

Tumpach AB, Drira H, Daoudi M, Srivastava A (2015) Gauge invariant framework for shape analysis of surfaces. *IEEE*

transactions on pattern analysis and machine intelligence 38(1):46–59

Turaga PK, Srivastava A (2016) Riemannian computing in computer vision. Springer

Vaillant M, Glaunès J (2005) Surface matching via currents. In: Biennial International Conference on Information Processing in Medical Imaging, Springer, pp 381–392

Willmore TJ (1993) Riemannian geometry. Oxford University Press

Younes L (1998) Computable elastic distances between shapes. SIAM J Appl Math 58(2):565–586

Younes L (2010) Shapes and diffeomorphisms, vol 171. Springer Science & Business Media

Younes L, Michor PW, Shah J, Mumford D (2008) A metric on shape space with explicit geodesics. Atti Accad Naz Lincei Rend Lincei Mat Appl 19(1):25–57