





# Divergence-Free Shape Correspondence by Deformation

M. Eisenberger , Z. Löhner  and D. Cremers 

Technical University of Munich, Germany



**Figure 1:** Given two input shapes we propose to morph the source shape along a divergence-free deformation field in order align it with the target. (Left) Example of a deformation field in 3D. (Right) Example of the results of our framework. We alternate between optimizing for the deformation field and calculating correspondences and therefore generate highly accurate correspondences (color coded) as well a sequence of natural intermediate shapes as a by-product (white). The translation is only added for visibility.

## Abstract

We present a novel approach for solving the correspondence problem between a given pair of input shapes with non-rigid, nearly isometric pose difference. Our method alternates between calculating a deformation field and a sparse correspondence. The deformation field is constructed with a low rank Fourier basis which allows for a compact representation. Furthermore, we restrict the deformation fields to be divergence-free which makes our morphings volume preserving. This can be used to extract a correspondence between the inputs by deforming one of them along the deformation field using a second order Runge-Kutta method and resulting in an alignment of the inputs. The advantages of using our basis are that there is no need to discretize the embedding space and the deformation is volume preserving. The optimization of the deformation field is done efficiently using only a subsampling of the original shapes but the correspondence can be extracted for any mesh resolution with close to linear increase in runtime. We show 3D correspondence results on several known data sets and examples of natural intermediate shape sequences that appear as a by-product of our method.

## 1. Introduction

Handling non-rigid, nearly isometric deformations of 3D shapes is at the heart of numerous problems in computer vision and graphics. Applications range from shape comparison, information and style transfer to the automatic generation of new, meaningful shapes. In comparison to rigid shape registration, the complexity increases significantly in the presence of non-rigid deformations. Many methods rely on purely intrinsic or very local measures because these are robust under extreme extrinsic changes [OBS\*12; ADK16]. While this helps to reduce the complexity of the correspondence problem, it also often leads to artifacts in the matching

coming from intrinsic symmetries, which are indistinguishable in the intrinsic view, or areas with indiscriminative features.

A different line of work aims at deforming shapes by directly manipulating their geometry in the embedding space [vFTS06; MS10; MZT\*14]. Having an explicit notion of extrinsic deformations yields more regular, continuous matchings which is relevant in many applications. In particular, this approach allows for the creation of new, intermediate versions of the input shapes. On the other hand, these methods are in general more prone to get stuck in local minima and therefore dependent on a good initial alignment of the inputs. Unfortunately, many extrinsic matching methods use linear

mappings to model surface deformations [MS10; MZT\*14]. While this is feasible for small changes, it is often not compatible with how objects deform in the real world. On the other hand, finding a physically correct morphing between two shapes is highly complex and computationally intense, even when the perfect correspondence or prior knowledge about the input is given [WBR11; GCLX17]. In this paper, we propose a more plausible morphing model that takes into account volume-preservation during the entire deformation. This is possible by modeling volume-preservation through zero divergence in a deformation field. This property makes our intermediate shapes more natural and our results are less likely to end up in a local minimum than with a linear mapping. In our method, the deformation field is represented in a spatially continuous, coarse-to-fine basis which allows for an efficient optimization. Moreover, we can apply the final deformation to shapes of arbitrary resolution with a minimal increase in complexity.

## 2. Related Work

### 2.1. Shape Registration and Matching

Much work has been done in the direction of shape registration and matching and we would like to point the interested reader to in-depth surveys of these topics for an overview [vKZHC11; SMFF07; TCL\*13]. Here we will focus on work that is directly related to our approach.

One recent line of work in shape matching is based on spectral decomposition of the surface Laplace-Beltrami operator [DK10]. This is popular because it reduces the dimensionality of the problem from the number of vertices to the number of basis functions chosen [OBS\*12]. Nevertheless, extracting the correspondence from the low dimensional representation is still a complex problem and often retrieved solutions are noisy or hard to compute [RMC15]. One major problem with purely spectral approaches is that intrinsic symmetries can not be distinguished, [RPWO18] being one of few exceptions. We also use a spectral approach but, instead of a basis for functions on the surface, we represent deformation fields in the embedding space using the eigenfunctions of the standard Laplacian. Among other things the embedding space allows us to distinguish between intrinsically symmetric but opposite points.

Methods based on Multi-Dimensional Scaling find correspondences by reembedding and then aligning shapes in a (possibly smaller) embedding space where the complexity is reduced [BBK06; ADK16]. [CK15] calculate a robust non-rigid registration based on Markov random fields but cannot retrieve a continuous deformation which we do. In [MS10] and [MZT\*14] the authors address the non-rigid registration problem by modeling one point cloud as a Gaussian mixture model, similar to our method. Moreover, they also determine the correspondences and point mappings in an alternating manner using a expectation maximization algorithm. This work is strongly related to our framework but no intermediate deformation is modeled. There also exist extensions of this method which additionally include descriptor values [MZY16; MJLL17]. [HAWG08] achieve accurate non-rigid alignments but rely on good initial correspondence and expensive geodesic distance computation to find these.

### 2.2. Deformation Fields

Deformation fields have a long history in image registration. One of the first approaches in that direction is the LDDMM framework [BMTY05]. Ashburner and colleagues made use of deformation fields for autonomous shape morphing [Ash07]. They consider temporally constant deformation fields offering limited flexibility to capture more complex deformations. Solving for a space and time dependent deformation field is a highly underdetermined problem. A remedy for this issue is provided by the geodesic shooting approach advocated by [MTY06] which only estimates the initial velocity field for each pixel and then how the velocity has to propagate in the image domain in order to preserve the kinetic energy and the momentum of the whole system. Further improvements of this framework were proposed in subsequent work, including a Gauss-Newton approach [AF11] and a particularly efficient adjoint calculation [VRR12].

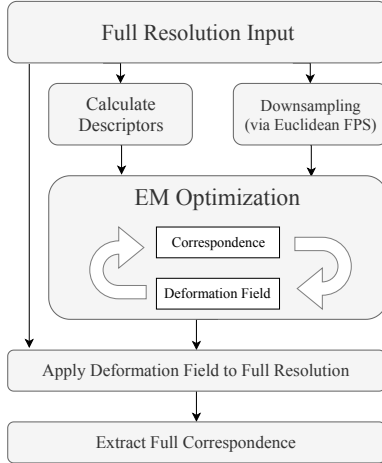
Closely related to our work is [vFTS06] in which the authors also model volume preserving shape deformations using divergence-free vector fields. Here, deformation fields are constructed from hand crafted templates which are meant to be used as interactive shape transformation tools whereas our method is fully automated. As in our work, in [AOW\*08] the deformations are based on a sub-sampling of the input shapes and can be efficiently applied to the full resolution but the correspondence is assumed to be given.

Probabilistic interpretations of deformation fields are a popular formulation. Such a model for image registration and 2D shape registration with a Gaussian process modeling of the correspondence mapping is proposed in [ALV08]. Further work [LJGV16; DGL\*17] specified how one can extend this approach to Gaussian processes on the surface of a three dimensional shape. The authors in [BHB00], [THB08], [ALV08] and [PDS\*09] also model non-rigid transformations using a PCA type representation of permitted motions. Analogously, [MS10] and [MZT\*14] pursue a reproducing kernel Hilbert space approach to model the vector field interpolation. However, for all these references the respective vector fields are not defined on the whole embedding space surrounding the shapes but rather only at the elements of the considered point clouds and they do not admit an interpretation as a deformation field which makes it harder to impose global properties, e.g. volume-preservation.

Another classical approach to shape deformation is based on a rotation invariant representation of triangle meshes [LSLC05]. In [ZSC\*08] this deformation model is used to compute a sparse set of correspondences but this method is hard to scale to high resolutions.

## 3. Contribution

We introduce a mathematical framework which solves the correspondence problem on two shapes with approximately the same volume. For this purpose, we propose to alternately estimate the correspondences and a smooth 3D deformation field aligning the two input shapes. Our shape morphing model solves an initial value problem to shift the first shape along this deformation field. Numerically, this differential equation is integrated using a second order Runge-Kutta scheme. Our framework allows us to incorporate



**Figure 2:** Overview over our complete pipeline.

physical assumptions about the deformations by directly building them into our model. We suggest to impose volume preservation by enforcing the deformation fields to have zero divergence. More specifically, we define a coarse-to-fine basis representation of these vector fields where each basis function is divergence-free. This allows us to reduce the complexity by optimizing only over the most significant coefficients. We use an expectation-maximization approach to simultaneously compute a subset of the unknown point-to-point correspondences and the optimal deformation field coefficients. A schematic diagram of the complete pipeline can be found in Figure 2. We demonstrate that the proposed framework can be used to solve for correspondences which are on par with state-of-the-art methods. Moreover, our method can produce a sequence of reasonable intermediate shapes between the inputs as a by-product. Both can be scaled up to arbitrary resolution without a significant increase in complexity which we demonstrate on a dataset of real scans with over 100k vertices.

## 4. Problem Formulation

In the following, we define the problem we want to solve and the mathematical background we use in later sections. In general we consider two point clouds  $\mathcal{X} = \{x_1, \dots, x_N\} \subset \Omega$  and  $\mathcal{Y} = \{y_1, \dots, y_M\} \subset \Omega$  contained in a compact domain  $\Omega \subset \mathbb{R}^D$ . In practice we choose  $\Omega = [0, 1]^D$ . These points  $x_n$  and  $y_m$  are samples from the surface of two similar  $(D-1)$ -dimensional Riemannian manifolds embedded in  $\mathbb{R}^D$ . Our method aims at aligning the point clouds  $\mathcal{X}$  and  $\mathcal{Y}$  in a meaningful manner. In particular, we are looking for a mapping  $f: \mathcal{X} \rightarrow \Omega$  which provides the coordinates for a new embedding of each point on  $\mathcal{X}$ . In the end,  $f(\mathcal{X})$  should be well aligned with  $\mathcal{Y}$ .

### 4.1. Deformation field shape morphing

We propose to model the shape morphing  $f: \mathcal{X} \rightarrow \Omega$  using the following initial value problem:

$$\begin{cases} \dot{x}(t) = v(x(t)), \\ x(0) = x_{\text{init}}. \end{cases} \quad (1)$$

In this context,  $v: \Omega \rightarrow \mathbb{R}^D$  is some fixed deformation field shifting any point  $x_{\text{init}} \in \Omega$  over time. If we solve this differential equation until some fixed time  $t_{\text{eval}}$ , we get the flow  $\varphi: [0, t_{\text{eval}}] \times \Omega \rightarrow \Omega$  of Equation (1). The flow  $\varphi$  morphs the space  $\Omega$  over time, it maps any input point  $x_{\text{init}}$  to its destination  $\varphi(t, x_{\text{init}})$  at time  $t \in [0, 1]$ . Applying Equation (1) to all points  $x_{\text{init}} := x_n \in \mathcal{X}$  yields a morphing model for the source shape  $\mathcal{X}$ :

$$f(x_n) := f_n := \varphi(t_{\text{eval}}, x_n). \quad (2)$$

In order to make those shape deformations more plausible, we require them to be smooth in space and in time. For this purpose, we assume that the deformation fields  $v \in C^\infty(\Omega, \mathbb{R}^D)$  which, according to the Picard-Lindelöf Theorem, yields smooth point trajectories  $x(\cdot) := \varphi(\cdot, x_{\text{init}}) \in C^\infty([0, 1], \Omega)$ , see [Tes12, Lemma 2.3, Theorem 2.5]. For convenience we choose  $t_{\text{eval}} = 1$  in our experiments.

Our morphing model computes natural shape deformations which can be transformed into correspondences through nearest neighbor search (See Section 5.3). Due to the time dependency of the flow, we additionally get intermediate poses of the input shape at times  $t \in (0, 1)$  which constitute the underlying transformation. Those are typically more meaningful than naive approaches like linear interpolation between the points. We believe that having a continuous correspondence and a natural deformation are inherently connected and solving for both simultaneously improves the results considerably.

### 4.2. Divergence-free deformations

One advantage of our morphing model (1) is that it allows us to incorporate assumptions about the deformation fields  $v$  into our model. In our framework we restrict these velocity fields to be divergence-free, an assumption that is commonly used in mathematical modeling of incompressible fluids [CM93]:

$$\nabla \cdot v = 0. \quad (3)$$

A well known consequence of this local property is that it yields volume preservation over time for any subpart  $U \subset \Omega$  of the embedding space. In particular, we can consider the set of solutions of Equation (1):

$$U(t) := \{\varphi(t, x_{\text{init}}) \in \Omega \mid x_{\text{init}} \in U\}. \quad (4)$$

Then the assumption in Equation (3) yields that each morphed set  $U(t)$  has the same volume as  $U$  [Tes12, Lemma 8.8]. Therefore, each subvolume of the input shape  $\mathcal{X}$ , as well as of the embedding space, is preserved at any given time. Notice that this property is stronger than global volume preservation of the interior of  $\mathcal{X}$  only. In general, two very differently shaped objects can have the same volume. However, for our method the volume of all, potentially very small, subparts is preserved. In our experiments, we found that this is a reasonable assumption for real world deformations and it provides a good regularization of our morphing model (1).

### 4.3. Helmholtz decomposition

Helmholtz's theorem [Ari62] implies that any sufficiently smooth vector field on the compact domain  $\Omega$  can be decomposed into

the sum of a curl-free, a divergence-free and a harmonic component. It furthermore provides us with an explicit construction of the divergence-free component that we are interested in:

$$v := \nabla \times \Phi. \quad (5)$$

In this context,  $\Phi : \Omega \rightarrow \mathbb{R}^D$  is a potential function and  $\nabla \times \cdot$  is the curl operator. Indeed, [CNT15, Lemma 2.2] shows that such a  $\Phi$  exists for any divergence-free,  $C^\infty$  vector field  $v : \Omega \rightarrow \mathbb{R}^D$  with no outflow at the boundary:

$$\langle v, n \rangle = 0 \text{ on } \partial\Omega. \quad (6)$$

Furthermore, for a given  $\Phi$  we always get a divergence-free vector field  $v$  as a basic property of the curl operator:

$$\nabla \cdot (\nabla \times \Phi) = 0. \quad (7)$$

To further restrict the space of admissible deformation fields, we additionally require the potential functions to admit Dirichlet boundary conditions  $\Phi|_{\partial\Omega} = 0$ . This guarantees that the potential functions are tangential to the outer normals at  $\partial\Omega$  which is a necessary condition in the existence proof [CNT15, Lemma 2.2]. Moreover, we are only interested in a high expressibility in the interior of  $\Omega$  and choosing Dirichlet boundary conditions makes the representation of our deformation fields even more compact. Intuitively, it guarantees that for the resulting deformation fields  $v$  there is no flow in and out of the domain  $\Omega$  (see Equation (6)). In the case of  $D = 3$  spatial dimensions the construction of  $v$  in (5) admits the following form:

$$v = \begin{pmatrix} \partial_2\Phi_3 - \partial_3\Phi_2 \\ \partial_3\Phi_1 - \partial_1\Phi_3 \\ \partial_1\Phi_2 - \partial_2\Phi_1 \end{pmatrix} = \begin{pmatrix} 0 \\ \partial_3\Phi_1 \\ -\partial_2\Phi_1 \end{pmatrix} + \begin{pmatrix} -\partial_3\Phi_2 \\ 0 \\ \partial_1\Phi_2 \end{pmatrix} + \begin{pmatrix} \partial_2\Phi_3 \\ -\partial_1\Phi_3 \\ 0 \end{pmatrix}. \quad (8)$$

**Remark.** The harmonic component in the Helmholtz decomposition corresponds to global translations of the input shape  $\mathcal{X}$  but we refrain from including them in our framework. For once, we would like the flow  $\phi : [0, 1] \times \Omega \rightarrow \Omega$  to map all points  $x_n \in \Omega$  back to the same domain. Furthermore, modeling global translations is not necessary because we shift the input shapes a priori such that their empirical mean corresponds to the center of  $\Omega$ .

## 5. Method

In the following, we outline the core components of our method. First, we construct a coarse-to-fine deformation field basis with certain built-in properties like volume preservation (Section 5.1). Then, we show how to integrate the initial value problem of Equation (1) (Section 5.2). Finally, we provide details about our expectation maximization algorithm (Section 5.3) where we simultaneously optimize for the unknown correspondences and an appropriate deformation field. We refer the reader to the supplementary material for full details. Regarding relevant applications, we will mainly restrict ourselves to the case of 2D shapes embedded in  $\mathbb{R}^3$ . However, extensions to higher dimensions  $D > 3$  or  $D = 2$  are straightforward.

### 5.1. Spatial representation

Standard discretizations of vector fields  $v$  using voxel grids have cubic complexity which makes them too costly for any reasonable resolution. To get a more compact representation, we introduce a low rank basis  $\{v_1, \dots, v_K\}$  of spatially dense, divergence-free deformation fields. The number of basis functions can be adjusted for either speed or expressiveness. Without loss of generality we set the domain to a  $D$ -dimensional cube  $\Omega := [0, 1]^D$ . In practice, we then translate and scale any shape to generously fit inside. We begin with defining a basis for the potential fields  $\Phi$ . For this purpose, consider the eigenfunctions  $\{\phi_1, \phi_2, \dots\}$  and eigenvalues  $\{\lambda_1, \lambda_2, \dots\}$  of the scalar Laplacian  $\Delta$  on  $\Omega$ :

$$\Delta\phi_k = \lambda_k^\Delta \phi_k. \quad (9)$$

This basis of eigenfunctions  $\{\phi_1, \phi_2, \dots\}$  is ordered with descending eigenvalues  $0 \geq \lambda_1^\Delta \geq \lambda_2^\Delta \geq \dots$ . Furthermore, we require the potential fields to admit Dirichlet boundary conditions  $\Phi|_{\partial\Omega} = 0$ . These  $\phi_k$  can be computed analytically, they are exactly the sine elements of the Fourier basis:

$$\mathcal{B}_\phi = \left\{ \phi : [0, 1]^D \rightarrow \mathbb{R}, x \mapsto \prod_{d=1}^D \frac{1}{2} \sin(x_d \pi j_d) \mid j \in \mathbb{N}^D \right\}. \quad (10)$$


The set  $\mathcal{B}_\phi = \{\phi_1, \phi_2, \dots\}$  is ordered by ascending Dirichlet energy of the  $\phi_k$ . These  $\phi_k$  form an orthonormal basis wrt. the  $\langle \cdot, \cdot \rangle_{L^2(\Omega)}$  inner product for scalar functions on  $\Omega$ . We can now use  $\mathcal{B}_\phi$  to construct a basis for the deformation fields  $\mathcal{B}_v$  according to Equation (5). Note that the basis  $\mathcal{B}_\phi$  consists of scalar functions while the potential functions  $\Phi : \Omega \rightarrow \mathbb{R}^D$  are vector valued. However, due to the linearity of the curl  $\nabla \times \cdot$  we obtain a basis by using (5) one entry at a time. For  $D = 3$  this can be done as follows:

$$\mathcal{B}_v = \bigcup_{k=1}^{\infty} \left\{ \nabla \times \begin{pmatrix} \phi_k \\ 0 \\ 0 \end{pmatrix}, \nabla \times \begin{pmatrix} 0 \\ \phi_k \\ 0 \end{pmatrix}, \nabla \times \begin{pmatrix} 0 \\ 0 \\ \phi_k \end{pmatrix} \right\} = \bigcup_{k=1}^{\infty} \left\{ \begin{pmatrix} 0 \\ \partial_3\phi_k \\ -\partial_2\phi_k \end{pmatrix}, \begin{pmatrix} -\partial_3\phi_k \\ 0 \\ \partial_1\phi_k \end{pmatrix}, \begin{pmatrix} \partial_2\phi_k \\ -\partial_1\phi_k \\ 0 \end{pmatrix} \right\}. \quad (11)$$

We get three deformation basis functions for each  $\phi_k$  in (10). Analogously to the potential fields, the basis elements  $\mathcal{B}_v = \{v_1, v_2, \dots\}$  are again sorted according to the eigenvalues  $\lambda_k^\Delta$  of the corresponding  $\phi_k$  in descending order. Note that there are in general multiple basis functions  $v_k$  for each eigenvalue  $\lambda_k^\Delta$ . Overall, we obtain arbitrary deformation fields  $v$  as the linear combination of the first  $K$  basis elements  $v_k$  with some coefficients  $a_k$ :

$$v(x) = \sum_{k=1}^K v_k(x) a_k. \quad (12)$$

**Remark** One aspect we would like to discuss in this context is our choice of domain  $\Omega = [0, 1]^D$ . The first basis function  $v_1$  in Figure 3 is equivalent up to first order to a rotation around the  $x_3$  axis. This especially holds near the center of the domain  $\Omega$  and deteriorates at its boundary  $\partial\Omega$ . Those considerations raise the question whether a cubic domain  $\Omega$  is the best choice for our purposes. Following the work in [ZB07; ZB08] we could pursue our approach in a spherical



**Figure 3:** Cross section of some deformation field basis functions  $v_k : \Omega \rightarrow \mathbb{R}^3$  at  $x_3 = 0.5$ . Notice the low frequency structures for low  $k$  and increasing frequencies with higher indices. Furthermore, one can see that our deformation fields have no flow in and out of the domain  $\Omega$  at the boundary.

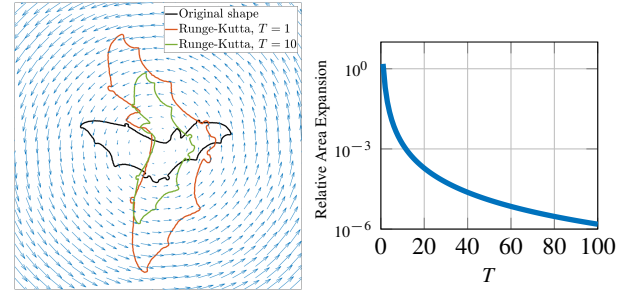
domain. This would lead to more complex basis functions  $v_k$  but the first three eigenfunctions would span the space of rotations without undesirable artifacts at the boundaries of the domain. Although this would be a nice theoretical property, we refrain from using these basis functions here due their complex structure.

## 5.2. Temporal discretization

In order to evaluate the correspondence mapping  $f$  in (2) we have to solve the initial value problem (1) with a numerical integration scheme. The simplest choice in this context is the explicit Euler method. However, we decided to use a second order Runge-Kutta method [GH10, Ch. 9] because it has a significantly higher accuracy and therefore allows for a coarser time discretization. We subdivide the time domain in an equidistant grid with  $T \in \mathbb{N}$  intervals and set the step size  $h = \frac{1}{T}$ . This yields an explicit iteration scheme:

$$\begin{cases} x_n^{(0)} := x_n. \\ x_n^{(t+1)} := x_n^{(t)} + hv \left( x_n^{(t)} + \frac{h}{2} v(x_n^{(t)}) \right). \\ f_n := x_n^{(T)}. \end{cases} \quad (13)$$

We typically choose  $T \in \{1, \dots, 100\}$  in our experiments. In general, we have to make a trade off between runtime and accuracy



**Figure 4:** Area expansion with different step sizes using the Runge-Kutta integration. Left: Rotation around 90 degrees on a bat shape of the MPEG-7 dataset [Ral] (black). If executed in one step ( $T = 1$ ) the shape expands (red) whereas for ten steps  $T = 10$  the area of the interior stays nearly the same (green). Right: Relative area expansion when performing the same deformation with an increasing number of steps  $T$ .


when selecting an appropriate number of steps  $T$ . If we choose  $T$  too small, we lose some key properties of our framework like the volume preservation. This effect is illustrated in Figure 4 for the 2D shape of a bat transformed by a 90 degree rotation around the center. Note that the deformation field corresponding to this transformation is only approximately contained in our framework due to our choice of domain and boundary conditions, see discussion in the previous subsection. If we choose too few time steps  $T$ , the shape shifts outward and the area expands. On the other hand, this effect becomes insignificantly small if we choose  $T \geq 10$ .

## 5.3. Optimization

In the previous sections we derived a coherent description of shape morphing using volume preserving deformation fields. We can now use this framework to construct an algorithm that matches two given point clouds  $\mathcal{X}$  and  $\mathcal{Y}$  by calculating a volume preserving deformation field between them. In order to do that we simultaneously optimize for the deformation field coefficients  $a$  and the unknown correspondences.

Similar to [MS10] and [MZT\*14] we approach shape registration in a probabilistic manner. We interpret the point cloud  $\mathcal{X}$  as a Gaussian mixture model with the means located at the shifted points  $f_n = x_n^{(T)}$  and the covariance  $\sigma^2 I_D \in \mathbb{R}^{D \times D}$  for some  $\sigma > 0$ . This enables us to simultaneously determine the deformation field coefficients  $a \in \mathbb{R}^K$  and the correspondences  $W \in [0, 1]^{N \times M}$  by applying an expectation maximization approach. Expectation maximization alternates between optimizing the deformation field coefficients and the correspondence while assuming the other to be fixed. The full derivation of each step, as well as additional implementation details, can be found in the supplementary material.

**Expectation step** The expectation step calculates correspondences for a fixed deformation. We represent the correspondences between the morphed  $f(\mathcal{X}) = \{f_1, \dots, f_N\}$  and the reference pointcloud  $\mathcal{Y} = \{y_1, \dots, y_M\}$  as soft correspondence matrices  $W \in [0, 1]^{N \times M}$  which arise naturally from the Gaussian mixture model assumption. High values of  $W_{nm} \approx 1$  indicate a high correspondence probability



**Figure 5:** Quantitative evaluation using the Princeton benchmark protocol on the TOSCA data set [BBK08] (left), the SCAPE data set [ASK\*05] (middle), and the high-resolution TOPKIDS [LRB\*16] (right). On TOSCA and SCAPE we compare against Spectral Generalized Multi-Dimensional Scaling (SGMDS) [ADK16], Functional Maps [OBS\*12], Blended Intrinsic Maps (BIM) [KLF11], Möbius Voting [LF09], Coherent Point Drift (CPD) [MS10] and Kernel Matching [VLB\*17]. On TOPKIDS we compare against the competitors of the original paper (Isometric Embedding via Expectation-Maximization (IE-EM) [SY12], Green’s Function Embedding Alignment (GFE) [BDK17], Random Forests (RF) [RRW\*14], Fully Spectral Partial Matching (FSPM) [LRBB17], Partial Functional Maps (PFM) [RCB\*17] and Kernel Matching (KM) [VLB\*17]. Both the TOSCA as well as the TOPKIDS dataset contain cases which are critical for our method but our results are still on a par with state-of-the-art. See Section 6.1 for details. On the TOSCA data set we additionally evaluate our method without using features. The drop in performance shows that these are crucial to avoid unwanted optima.

for the point pair  $(x_n, y_m)$  while values close to zero indicate low probability. The expectation maximization framework then yields an explicit update rule for  $W$  given the deformation coefficients  $a$ :

$$W_{nm} := \frac{\exp\left(-\frac{1}{2\sigma^2}d_{nm}^2\right)}{(2\pi\sigma^2)^{\frac{D}{2}} + \sum_{\tilde{n}=1}^N \exp\left(-\frac{1}{2\sigma^2}d_{\tilde{n}m}^2\right)}. \quad (14)$$

For a derivation of this formula, see the supplementary material. Intuitively,  $W_{nm}$  describes the value of a Gaussian with center  $f_n$  and variance  $\sigma$  at point  $y_m$ . Similar to [MS10], the normalization factor in the denominator comes from the mixture model assumption combined with an explicit modeling of outliers. In order to prevent our method from getting stuck in incorrect local optima, we include SHOT descriptor [TSS10] with standard parameters from the authors’ implementation. We combine them with Euclidean distances to define a metric for pairs of points  $x_n$  and  $y_m$ :

$$d_{nm}^2 := \|y_m - f_n\|_2^2 + \bar{d} \|\text{SHOT}(x_n) - \text{SHOT}(y_m)\|_2^2. \quad (15)$$

We introduce the factor  $\bar{d} \geq 0$  to ensure that both metrics have a comparable scaling, in particular we require both summands to have the same mean value for all point pairs  $\mathcal{X}$  and  $\mathcal{Y}$ . Note that we use descriptor values  $\text{SHOT}(x_n)$  on the original shape  $\mathcal{X}$  instead of the morphed shape  $f(\mathcal{X})$  in order to not recompute them at every iteration.

**Maximization step** The maximization step updates the deformation field for given soft correspondences  $W$ . Intuitively, we are looking for the deformation field coefficients  $a$  that best align points with high correspondence probability  $W_{nm}$ . For this purpose, we interpret the coefficients  $a = (a_1, \dots, a_K)^\top$  as random variables with

a normal distribution  $a \sim \mathcal{N}(0, L)$ , where  $L := \text{diag}(\lambda_1, \dots, \lambda_K)$ . If we compute the pushforward of this Gaussian according to Equation (12), we get a prior distribution of deformation fields  $v$ . The weights  $\lambda_k$  are constructed from the eigenvalues  $\lambda_k^\Delta$  as follows:

$$\lambda_k := (-\lambda_k^\Delta)^{-\frac{D}{2}} = \left(\pi^2 \sum_{d=1}^D J_d^2\right)^{-\frac{D}{2}}. \quad (16)$$


The mathematical background of this choice for the weights  $\lambda_k$  is provided by the Karhunen-Loève expansion [Sul15, Ch. 11] which is an extension of the principal component analysis (PCA) for function spaces, see the supplementary material for more details. Intuitively, this kind of weighting promotes a damping of the high frequency components and smoothness of the deformation field  $v$ . The maximization step optimizes the coefficients  $a$  for their posterior distribution given the current correspondences which describes how well the deformation field of  $a$  explains  $W$ . This results in the following energy for  $a$ :

$$E(a) := \frac{\sigma^2}{2} a^\top L^{-1} a + \sum_{m=1}^M \sum_{n=1}^N W_{nm} \rho(\|y_m - f_n\|_2). \quad (17)$$

This energy  $E$  is the sum of the negative log prior including the weights  $\lambda_k$  (left term) and the negative log likelihood (right term) of  $a$ . The function  $\rho: \mathbb{R} \rightarrow [0, \infty)$  is the Huber loss [Hub64] which helps to account for outliers and makes the deformation field estimation more robust:

$$\rho(r) = \begin{cases} \frac{1}{2}r^2 & |r| \leq r_0 \\ r_0|r| - \frac{1}{2}r_0^2 & \text{otherwise.} \end{cases} \quad (18)$$

In our experiments, we choose the outer slope as  $r_0 := 0.01$ . Furthermore, we apply a Gauss-Newton type approach to minimize



**Figure 6:** Examples of texture transfer done with our method. For each object the first image shows the source shape and texture, the second image the texture transferred with the ground-truth map and the third image the texture transferred with our correspondences. Our results are nearly identical to the ground-truth except for the dog which shows some artifacts on tail and chest.

the energy in (17). This results in an iterative method similar to the Levenberg-Marquardt algorithm [Lev44]. For this purpose, the residual term  $\|y_m - f_n\|_2$  is linearized in each iteration. This requires a differentiation of the Runge-Kutta scheme (13) wrt. the weights  $a$ , see the supplementary for an explicit formulation of the derivative  $\frac{d}{da}f_n$  and the Gauss-Newton update step for the energy in Equation (17).

To summarize, our method alternates between computing the weights  $W^{(i)}$  according to (14) and performing one Gauss-Newton update step for (17) to obtain  $a^{(i)}$ . To initialize the algorithm we set the deformation field to zero  $a^{(0)} := 0$ .

## 6. Experiments

We evaluate our method for several applications to show that it is general and flexible. Although we handle shapes with up to 200k and more vertices, the computation of the deformation field is always done on a downsampled version of the inputs with 3000 vertices and then applied to the full resolution. We use Euclidean farthest point sampling. The downsampled shape should include points of all relevant large and fine scale structures in order for the deformation field to move these correctly but we found 3000 sufficient for our applications. As a preprocessing step we shift both inputs such that the mean of their vertex positions is in the center of the domain and align them using PCA. To avoid wrong alignments along the principle component axes we choose the orientation that minimizes Eq. (15). When averaging over all experiments presented here, our algorithm takes about 370 seconds to compute the deformation and correspondences for one pair of shapes. Due to our a priori downsampling the runtime is only linearly dependent on the number of vertices, see Section 6.4 for a discussion of this property. All experiments were performed with MATLAB on a system with an Intel Core i7-3770 CPU clocked at 3.40GHz, 32 GB RAM and a GeForce GTX TITAN X graphics card running a recent Linux distribution. In all our experiments we only use the

raw shape data and in particular do not need any ground truth information or user input.


### 6.1. Matching

We verify our method using the TOSCA [BBK08], SCAPE [ASK\*05] and high-resolution TOPKIDS [LRB\*16] data sets. All these shapes are synthetic and therefore the exact intraclass correspondences are known. TOSCA contains 76 triangular meshes with 8 classes of humans and animals, SCAPE consists of 72 poses of the same person and TOPKIDS contains 26 poses of the same person in which topological merging, as it might appear in real scanning, is imitated.

We set the hyperparameters  $\sigma^2 := 0.01$ ,  $T := 20$  and choose  $K = 3000$  basis functions for the deformation field. Because  $W^{(i)}$  only contains 3000 correspondences we perform a nearest-neighbor search with respect to the metric in Eq. (15) to obtain a dense mapping. The evaluation is done with the Princeton benchmark protocol [KLF11]. Given the ground-truth match  $(x, y^*) \in \mathcal{X} \times \mathcal{Y}$ , the error of the calculated match  $(x, y)$  is given by the geodesic distance between  $y$  and  $y^*$  normalized by the diameter of  $\mathcal{Y}$ :

$$\epsilon(x) = \frac{d_{\mathcal{Y}}^{\text{Geo}}(y, y^*)}{\sqrt{\text{area}(\mathcal{Y})}}$$

We plot cumulative curves showing the percentages of matches that are below an increasing threshold. As zero is the value for ground-truth matches, the ideal curve would be constant at 100. See Figure 5 for our results and Figure 6 for example matching results showing texture transfer. On SCAPE we are able to reach state-of-the-art results whereas on TOSCA the intrinsic Kernel Matching methods is slightly better. Our extrinsic approach makes self-touching poses more challenging and these cases occur fairly often in TOSCA. Although TOPKIDS is still synthetic the self-touching poses are actually merged in the geometry which makes it more challenging. On this dataset we are slightly better than Kernel



**Figure 7:** Example registrations from the FAUST scan data set. The surface color corresponds to the Euclidean surface distance between scan and registration. The scale of the scans is in real cm values and the same on all plots. We report the average and maximum error under each image. Many errors occur due to the SHOT descriptors being corrupted at holes and in noisy areas (e.g. the hands). Furthermore, in some case the assumption of exact volume preservation is too restrictive for real scans with noise and topological changes (see especially second to the right).

Matching (see Figure 5). See the supplementary material for an example.

To show the influence of features on the results we do an evaluation of our method without using features at any point during the optimization. Instead the distance of Equation (15) is replaced with the pure Euclidean distance between the coordinates. The result can be seen in Figure 5. The performance without features decreases substantially because the Euclidean distance is a weak indicator when large deformations take place. Therefore, our method gets stuck in local optima more often.

## 6.2. Registration

We apply our framework to the FAUST Scan dataset [BRLB14] which contains data from scans of real humans in different poses. Each of these shapes has approximately 200k vertices, they are sampled inconsistently and some of them are severely affected by scanning noise, holes and topological changes. We match the null shape of every person to its other poses. In Figure 7 we display the surface distance of the morphed shapes to the goal shape for some examples. We reach very tight alignments except in very challenging cases like topological changes. Furthermore, the scanned volume varies slightly even between different poses of the same humans which induces small errors in our method.

## 6.3. Effect of the basis size

In our evaluations we consistently use  $K = 3000$  deformation field basis functions. To justify this choice empirically, we compute the mean geodesic errors of each TOSCA pair for several basis sizes  $K \in \{1, \dots, 3000\}$ , see Figure 9. We observe that while the accuracy increases significantly for small  $K \leq 1000$ , after some point it starts plateauing. In our evaluations, we choose  $K = 3000$  because we aim for a high accuracy. However, for some applications


where runtime is more important than accuracy a smaller basis size  $K < 3000$  might be sufficient.

## 6.4. Runtime for high resolution

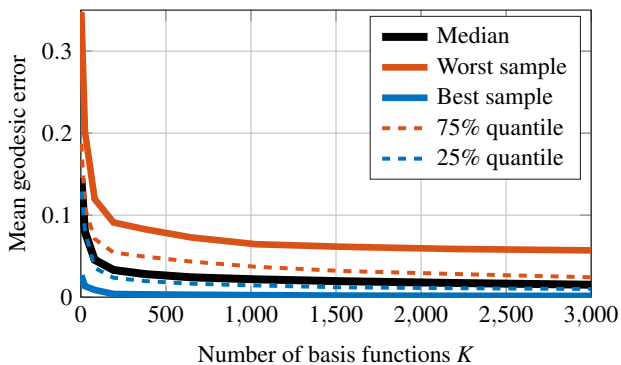
One major advantage of our method is that it is scalable to high resolution input shapes like those from FAUST because we optimize for the deformation field on downsampled shapes (3000 vertices). One point that we want to stress in this context is that this is not the same as computing matchings only on low resolution shapes. For many matching methods this scaling to the full resolution is challenging, most methods need to come up with a custom coarse-to-fine strategy. In general, it is not straightforward to extend a shape matching or deformation from a downsampled shape to the rest of the vertices. However, for our method this upscaling is trivial because the deformation field basis functions (11) are defined densely on the whole embedding space, therefore they can be evaluated anywhere in  $\Omega$ . This upsampling scales linearly in  $N$  because the Runge-Kutta method (13) is computed independently for all vertices  $x_n$ . See Figure 10 for an empirical verification of this property. Here, the runtime for the shape deformation is computed for various downsampled versions of one high resolution shape. To sum it up, the runtime for computing shape morphings is relatively low and increases only linearly in the number of vertices  $N$  which makes our method scalable for high resolution input shapes.

## 6.5. Shape Interpolation

**Interpolation** Our method morphs the input shape  $\mathcal{X}$  by solving the ODE (1) up to time  $t_{\text{eval}} = 1$ . If we now instead evaluate it at an intermediate time  $t \in (0, 1)$ , we get interpolated shapes as a byproduct of our matching pipeline. Just like the morphed shapes  $f(\mathcal{X})$  those intermediate shape morphings are smooth and volume preserving which makes them look natural. Three examples with




**Figure 8:** Three examples of shapes that are morphed into one another according to the initial value problem of Eq. (1). The centaur (a) and the human (b) are from the TOSCA [BBK08] and FAUST [BRLB14] dataset respectively. The armadillo (c) is from the AIM@SHAPE shape repository [AIM]. (b) is a scan of a real person and very high resolution (214k vertices). The source and target shape are shown in white and the interpolations at times  $t = 0.25, 0.5, 0.75$  in blue. The translation is not part of our deformation and was only introduced for clarity in the figures.



**Figure 9:** Dependency of the mean geodesic errors on TOSCA on-different basis sizes  $K \in \{1, \dots, 3000\}$ . In particular, we show the elements at the 0%, 25%, 50%, 75% and 100% quantile.

interpolated shapes are displayed in Figure 8 and videos can be found in the supplementary material.

**Extrapolation** Similarly to the idea of interpolating shapes as a byproduct of our method we can also use the computed deformation field  $v$  to solve the initial value problem (1) up to times  $t > 1$ . This results in extrapolated shapes, see Figure 11. In contrast to shape interpolation, extrapolation is a severely underdetermined task and it is hard to evaluate quantitatively. Nevertheless, we observed that for moderate time spans  $t \in [1, 1.5]$  our method produces reasonable results. In general, the morphing speed slows down at some point, especially when the shape is moving in previously unoccupied space. Intuitively, for the optimization there no incentive to impose any particular movement on these parts of the domain  $\Omega$ , if it is not relevant for the surface alignment. Still, our



**Figure 10:** Runtime of our method for the full resolution shape deformation for different number of vertices  $N \in \{3000, \dots, 100000\}$ . The full pipeline has two steps: (1) a fixed size optimization over 3000 vertices which takes around 360 seconds on average (blue dashed line), (2) applying the deformation field to the full resolution shape and extracting the correspondence for the full shape. The plot shows that our method scales linearly in the number of vertices and is therefore still feasible for very detailed shapes with over 100k vertices.

extrapolated shapes are visually appealing and not too severely affected by distortions.

## 7. Conclusion

We presented a novel extrinsic approach to shape matching. Besides computing dense surface correspondences, we also determine a smooth, volume preserving deformation field between the input shapes. Our morphing model shifts the source shape  $\mathcal{X}$  along this deformation field using a second order Runge-Kutta integration scheme in order to align it with the reference shape  $\mathcal{Y}$ . Additionally to aligning the inputs, this model can also be used to efficiently calculate plausible interpolated shapes.

Our method addresses the coupled problem of finding an unknown deformation and correspondence with an expectation maximization approach. Furthermore, we represent our morphing model with a low rank deformation field basis which reduces the degrees of freedom and makes the optimization problem well constrained. This then allows for a subsampling of the inputs which makes it computationally feasible, even for high resolution meshes, with only a linear increase in runtime. Quantitative evaluations for shape correspondence partly prove state-of-the-art performance of our method. Moreover, we show convincing examples of shape interpolation and extrapolation that arise naturally from our pipeline.

### 7.1. Limitations

Due to our choice of basis the deformation field is forced to be volume preserving. This makes sense in applications with the same object but prevents inter-class matchings - for example between two humans with different body shapes. The volume preservation property applies to every subregion of the domain  $\Omega$ , including the intermediate space between parts of the shape. Therefore, separating two touching parts (for example two hands) is in theory possi-



**Figure 11:** Example of an extrapolated shape from the KIDS dataset [RBW\*14]. It can be determined using the temporally fixed deformation field  $v$  for simulating the initial value problem (1) up to the time  $t = 1.3$ . Source and target shape are white, one interpolated shape is shown in blue and the extrapolation is pink. The deformation field is usually magnified in the area between the input shapes and fans out in several directions further away from the input shapes. Therefore, choosing a really high time does not lead to broken shapes but the movement slows down more and more until it basically stops.

ble but requires many high frequency deformation basis elements which would make the optimization costly.

The assumption of (1) being autonomous can be problematic if different parts of the shape move through the same region of the embedding space in a contradictory manner. One example for this is a hand closing to a fist. At first the index and middle finger occupy parts of the embedding space before the thumb moves in the same area but in a different direction. See Figure 12. A possible remedy for this problem is making the deformation fields time dependent.

Furthermore, since there is not one unique, volume preserving deformation between two shapes, our interpolation is not guaranteed to be as-rigid-as-possible which is a plausible assumption in many applications. If the displacement is spatially far, we might end up with squeezed intermediate states that are volume preserving but are affected by undesirable distortions. This is also visible in Figure 12 where the tip of the thumb becomes flat.

### 7.2. Future Work

Right now, our method will always find a solution that is globally volume preserving. This allows to find good deformations fields in the case of severe non-rigid deformations but is not applicable to partial data. In the future, we want to extend this method to work on real scans, for example from the Kinect, which naturally only show partial shapes. A promising approach for this is making the


deformation divergence-free in certain parts of the domain while still allowing volume changes in some places and imposing an as-killing-as-possible constraint in the optimization. Additionally, we want to construct an adaptive basis for this depending on the input shapes. This might also help with the separation of close parts and handling non volume preserving deformations like style or class changes. Furthermore, we only calculate one time independent field for the entire deformation which means mass at one spatial point always needs to move in the same direction at a later time step. This restricts the complexity of the deformations that our method can handle, especially for large-scale motions over a longer period of time. Future versions should allow more flexible types of deformation fields to extend it to a broader range of applications. We could for example associate different parts of the shape with different deformations fields or let them vary over time to address more difficult tasks.

### Acknowledgements

We thank Emanuele Rodolá, Nikolaus Demmel and Emanuel Laude for useful discussions. We gratefully acknowledge the support of the ERC Consolidator Grant "3D Reloaded".

### References

- [ADK16] AFLALO, YONATHAN, DUBROVINA, ANASTASIA, and KIMMEL, RON. "Spectral generalized multi-dimensional scaling". *IJCV* 118.3 (2016), 380–392 1, 2, 6.
- [AF11] ASHBURNER, JOHN and FRISTON, KARL J. "Diffeomorphic registration using geodesic shooting and Gauss–Newton optimisation". *NeuroImage* 55.3 (2011), 954–967 2.
- [AIM] AIM@SHAPE REPOSITORY. *AIM@SHAPE repository*. [http://visionair.ge.imati.cnr.it:8080/ontologies/shapes/viewgroup.jsp?id=657-Armadillo\\_19](http://visionair.ge.imati.cnr.it:8080/ontologies/shapes/viewgroup.jsp?id=657-Armadillo_19).



**Figure 12:** Example of a failure case. The source and target shapes are white, the interpolated shape at  $t = 0.5$  is blue and the resulting shape at  $t = 1$  is yellow. The yellow shape is supposed to be as close to the target as possible but the fingers are pushed away from each other instead. Here, the thumb and index finger are supposed to move in spatially very close areas. We are only calculating one deformation field for all time steps, therefore it is not clear for our method what motion to assign for contradicting motions. In order to resolve this problem, we need to either assign different motions to different subparts of the objects or make the deformation fields time dependent.

- [ALV08] ALBRECHT, THOMAS, LÜTHI, MARCEL, and VETTER, THOMAS. "A statistical deformation prior for non-rigid image and shape registration". *2008 IEEE Conference on Computer Vision and Pattern Recognition* (2008), 1–8 2.
- [AOW\*08] ADAMS, BART, OVSJANIKOV, MAKS, WAND, MICHAEL, et al. "Meshless Modeling of Deformable Shapes and their Motion". *Eurographics/SIGGRAPH Symposium on Computer Animation*. Ed. by GROSS, MARKUS and JAMES, DOUG. The Eurographics Association, 2008 2.
- [Ari62] ARIS, RUTHERFORD. *Vectors, tensors, and the basic equations of fluid mechanics*. Englewood Cliffs, N.J., Prentice-Hall, 1962 3.
- [Ash07] ASHBURNER, JOHN. "A fast diffeomorphic image registration algorithm". *NeuroImage* 38.1 (2007), 95–113 2.
- [ASK\*05] ANGUELOV, DRAGOMIR, SRINIVASAN, PRAVEEN, KOLLER, DAPHNE, et al. "SCAPE: shape completion and animation of people". *ACM transactions on graphics (TOG)*. Vol. 24. 3. ACM, 2005, 408–416 6, 7.
- [BBK06] BRONSTEIN, ALEXANDER M, BRONSTEIN, MICHAEL M, and KIMMEL, RON. "Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching". *PNAS* 103.5 (2006), 1168–1172 2.
- [BBK08] BRONSTEIN, ALEXANDER M, BRONSTEIN, MICHAEL M, and KIMMEL, RON. *Numerical geometry of non-rigid shapes*. [http://tosca.cs.technion.ac.il/book/resources\\_data.html](http://tosca.cs.technion.ac.il/book/resources_data.html). Springer, 2008 6, 7, 9.
- [BDK17] BURGHARD, OLIVER, DIECKMANN, ALEXANDER, and KLEIN, REINHARD. "Embedding Shapes with Green's functions for Global Shape Matching". *Computers & Graphics* 68C (2017), 1–10 6.
- [BHB00] BREGLER, CHRISTOPH, HERTZMANN, AARON, and BIEMANN, HENNING. "Recovering non-rigid 3D shape from image streams". *Proceedings IEEE Conference on Computer Vision and Pattern Recognition* 2 (2000), 690–696 2.
- [BMTY05] BEG, M FAISAL, MILLER, MICHAEL I, TROUVÉ, ALAIN, and YOUNES, LAURENT. "Computing large deformation metric mappings via geodesic flows of diffeomorphisms". *International journal of computer vision* 61.2 (2005), 139–157 2.
- [BRLB14] BOGO, FEDERICA, ROMERO, JAVIER, LOPER, MATTHEW, and BLACK, MICHAEL J. "FAUST: Dataset and evaluation for 3D mesh registration". *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. Piscataway, NJ, USA: IEEE, June 2014 8, 9.
- [CK15] CHEN, QIFENG and KOLTUN, VLADLEN. "Robust Nonrigid Registration by Convex Optimization". *IEEE International Conference on Computer Vision (ICCV)*. 2015, 2039–2047 2.
- [CM93] CHORIN, ALEXANDRE J and MARSDEN, JERROLD E. *A Mathematical Introduction to Fluid Mechanics*. Springer, 1993 3.
- [CNT15] CREUSÉ, EMMANUEL, NICAISE, SERGE, and TANG, ZUQI. "Helmholtz decomposition of vector fields with mixed boundary conditions and an application to a posteriori finite element error analysis of the Maxwell system". *Mathematical Methods in the Applied Sciences* 38.4 (2015), 738–750 4.
- [DGL\*17] DÖLZ, JÜRGEN, GERIG, T, LÜTHI, MARCEL, et al. "Efficient computation of low-rank Gaussian process models for surface and image registration". (2017) 2.
- [DK10] DUBROVINA, ANASTASIA and KIMMEL, RON. "Matching shapes by eigendecomposition of the Laplace-Beltrami operator". 2010 2.
- [GCLX17] GAO, LIN, CHEN, SHU-YU, LAI, YU-KUN, and XIA, SHIHONG. "Data-Driven Shape Interpolation and Morphing Editing". *Comput. Graph. Forum* 36 (2017), 19–31 2.
- [GH10] GRIFFITHS, DAVID and HIGHAM, DESMOND J. *Numerical Methods for Ordinary Differential Equations. Initial Value Problems*. Springer, 2010 5.
- [HAWG08] HUANG, QI-XING, ADAMS, BART, WICKE, MARTIN, and GUIBAS, LEONIDAS. "Non-Rigid Registration Under Isometric Deformations". *Computer Graphics Forum* 27 (2008), Number 5 2.

- [Hub64] HUBER, PETER J. “Robust Estimation of a Location Parameter”. *Annals of Statistics* 53.1 (1964), 73–101 6.
- [KLF11] KIM, VLADIMIR G., LIPMAN, YARON, and FUNKHOUSER, THOMAS A. “Blended intrinsic maps”. *Trans. Graphics* 30.4 (2011) 6, 7.
- [Lev44] LEVENBERG, KENNETH. “A method for the solution of certain non-linear problems in least squares”. *Quarterly of Applied Mathematics* 2.2 (1944), 164–168 7.
- [LF09] LIPMAN, YARON and FUNKHOUSER, THOMAS. “Möbius voting for surface correspondence”. *Trans. Graphics*. Vol. 28. 3. 2009, 72 6.
- [LJGV16] LÜTHI, MARCEL, JUD, CHRISTOPH, GERIG, THOMAS, and VETTER, THOMAS. “Gaussian Process Morphable Models”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 99 (2016) 2.
- [LRB\*16] LÄHNER, Z., RODOLA, E., BRONSTEIN, M. M., et al. “SHREC’16: Matching of Deformable Shapes with Topological Noise”. *Proc. of Eurographics Workshop on 3D Object Retrieval (3DOR)*. May 2016 6, 7.
- [LRBB17] LITANY, OR, RODOLÀ, EMANUELE, BRONSTEIN, ALEX, and BRONSTEIN, MICHAEL. “Fully Spectral Partial Shape Matching”. *Computer Graphics Forum* 36.2 (2017), 1681–1707 6.
- [LSLC05] LIPMAN, YARON, SORKINE, OLGA, LEVIN, DAVID, and COHEN-OR, DANIEL. “Linear rotation-invariant coordinates for meshes”. *ACM Transactions on Graphics (TOG)*. Vol. 24. 3. ACM. 2005, 479–487 2.
- [MJLL17] MA, JIAYI, JIANG, JUNJUN, LIU, CHENGYIN, and LI, YANSHENG. “Feature guided Gaussian mixture model with semi-supervised EM and local geometric constraint for retinal image registration”. *Information Sciences* 417 (2017), 128–142 2.
- [MS10] MYRONENKO, ANDRIY and SONG, XUBO. “Point Set Registration: Coherent Point Drift”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32.12 (2010), 2262–2275 1, 2, 5, 6.
- [MTY06] MILLER, MICHAEL I, TROUVÉ, ALAIN, and YOUNES, LAURENT. “Geodesic shooting for computational anatomy”. *Journal of mathematical imaging and vision* 24.2 (2006), 209–228 2.
- [MZT\*14] MA, JIAYI, ZHAO, JI, TIAN, JINWEN, et al. “Robust point matching via vector field consensus”. *IEEE Transactions on Image Processing* 23.4 (2014), 1706–1721 1, 2, 5.
- [MZY16] MA, JIAYI, ZHAO, JI, and YUILLE, ALAN L. “Non-rigid point set registration by preserving global and local structures”. *IEEE Transactions on image Processing* 25.1 (2016), 53–64 2.
- [OBS\*12] OVSJANIKOV, MAK, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. “Functional maps: a flexible representation of maps between shapes”. *ACM Transactions on Graphics (TOG)* 31.4 (2012), 30 1, 2, 6.
- [PDS\*09] PALADINI, MARCO, DEL BUE, ALESSIO, STOSIC, MARKO, et al. “Factorization for non-rigid and articulated structure using metric projections”. (June 2009), 2898–2905 2.
- [Ral] RALPH, RICHARD. *MPEG-7 Core Experiment CE-Shape-1 Test Set*. <http://www.dabi.temple.edu/~shape/MPEG7/dataset.html> 5.
- [RBW\*14] RODOLA, EMANUELE, BULO, SAMUEL, ROTA, WINDHEUSER, THOMAS, et al. “Dense Non-Rigid Shape Correspondence Using Random Forests”. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014 10.
- [RCB\*17] RODOLÀ, E., COSMO, L., BRONSTEIN, M. M., et al. “Partial Functional Correspondence”. *Computer Graphics Forum* 36.1 (2017), 222–236 6.
- [RMC15] RODOLÀ, EMANUELE, MOELLER, MICHAEL, and CREMERS, DANIEL. “Point-wise Map Recovery and Refinement from Functional Correspondence”. *Vision, Modeling and Visualization*. The Eurographics Association, 2015 2.
- [RPWO18] REN, JING, POULENARD, ADRIEN, WONKA, PETER, and OVSJANIKOV, MAK. “Continuous and Orientation-preserving Correspondences via Functional Maps”. *ACM Trans. Graph.* 37.6 (Nov. 2018), 248:1–248:16 2.
- [RRW\*14] RODOLÀ, EMANUELE, ROTA BULÒ, SAMUEL, WINDHEUSER, THOMAS, et al. “Dense non-rigid shape correspondence using random forests”. *Proc. CVPR*. 2014 6.
- [SMFF07] SALVI, JOAQUIM, MATABOSCH, CARLES, FOFI, DAVID, and FOREST, JOSEP. “A Review of Recent Range Image Registration Methods with Accuracy Evaluation”. *Image Vision Comput.* 25.5 (2007), 578–596 2.
- [Sul15] SULLIVAN, T J. *Introduction to Uncertainty Quantification*. Springer, 2015 6.
- [SY12] SAHILLIOĞLU, YUSUF and YEMEZ, YÜCEL. “Minimum-Distortion Isometric Shape Correspondence Using EM Algorithm”. *IEEE Trans. Pattern Anal. Mach. Intell.* 34.11 (2012), 2203–2215 6.
- [TCL\*13] TAM, GARY K. L., CHENG, ZHI-QUAN, LAI, YU-KUN, et al. “Registration of 3D Point Clouds and Meshes: A Survey from Rigid to Nonrigid”. *IEEE Transactions on Visualization and Computer Graphics* 19.7 (July 2013), 1199–1217 2.
- [Tes12] TESCHL, GERALD. *Ordinary Differential Equations and Dynamical Systems*. AMS, 2012 3.
- [THB08] TORRESANI, LORENZO, HERTZMANN, AARON, and BREGLER, CHRISTOPH. “Nonrigid structure-from-motion: estimating shape and motion with hierarchical Priors”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30.5 (2008), 878–892 2.
- [TSS10] TOMBARI, FEDERICO, SALTI, SAMUELE, and STEFANO, LUIGI DI. “Unique signatures of histograms for local surface description”. *In Proceedings ECCV* 16.9 (2010), 356–369 6.
- [vFTS06] VON FUNCK, WOLFRAM, THEISEL, HOLGER, and SEIDEL, HANS-PETER. “Vector field based shape deformations”. *ACM Transactions on Graphics (TOG)*. Vol. 25. 3. ACM. 2006, 1118–1125 1, 2.
- [vKZHC11] VAN KAICK, OLIVER, ZHANG, HAO, HAMARNEH, GHASAN, and COHEN-OR, DANIEL. “A Survey on Shape Correspondence”. *Computer Graphics Forum* 30.6 (2011), 1681–1707 2.
- [VLB\*17] VESTNER, MATTHIAS, LÄHNER, ZORAH, BOYARSKI, AMIT, et al. “Efficient Deformable Shape Correspondence via Kernel Matching”. *International Conference on 3D Vision (3DV)*. Oct. 2017 6.
- [VRR12] VIALARD, FRANÇOIS-XAVIER, RISSER, LAURENT, RUECKERT, DANIEL, and COTTER, COLIN J. “Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation”. *International Journal of Computer Vision* 97.2 (2012), 229–241 2.
- [WBRS11] WIRTH, BENEDIKT, BAR, LEAH, RUMPF, MARTIN, and SAPIRO, GUILLERMO. “A Continuum Mechanical Approach to Geodesics in Shape Space”. *International Journal of Computer Vision* 93.3 (July 2011), 293–318. ISSN: 1573-1405 2.
- [ZB07] ZHAO, CHUNYU and BURGE, JAMES H. “Orthonormal vector polynomials in a unit circle, Part I: basis set derived from gradients of Zernike polynomials”. *Opt. Express* 15.26 (2007), 18014–18024 4.
- [ZB08] ZHAO, CHUNYU and BURGE, JAMES H. “Orthonormal vector polynomials in a unit circle, Part II: completing the basis set”. *Opt. Express* 16.9 (2008), 6586–6591 4.
- [ZSC\*08] ZHANG, HAO, SHEFFER, ALLA, COHEN-OR, DANIEL, et al. “Deformation-driven shape correspondence”. *Computer Graphics Forum*. Vol. 27. 5. Wiley Online Library. 2008, 1431–1439 2.