

Spatial Maps: From Low Rank Spectral to Sparse Spatial Functional Representations

Andrea Gasparetto¹

andrea.gasparetto@unive.it

Luca Cosmo¹

luca.cosmo@unive.it

Emanuele Rodolà^{2,3}

rodola@di.uniroma1.it

Michael M. Bronstein^{3,4}

michael.bronstein@usi.ch

Andrea Torsello¹

torsello@unive.it

¹Università Ca' Foscari Venezia ²Sapienza University of Rome ³USI Lugano ⁴Intel

Abstract

Functional representation is a well-established approach to represent dense correspondences between deformable shapes. The approach provides an efficient low rank representation of a continuous mapping between two shapes, however under that framework the correspondences are only intrinsically captured, which implies that the induced map is not guaranteed to map the whole surface, much less to form a continuous mapping. In this work, we define a novel approach to the computation of a continuous bijective map between two surfaces moving from the low rank spectral representation to a sparse spatial representation. Key to this is the observation that continuity and smoothness of the optimal map induces structure both on the spectral and the spatial domain, the former providing effective low rank approximations, while the latter exhibiting strong sparsity and locality that can be used in the solution of large-scale problems. We cast our approach in terms of the functional transfer through a fuzzy map between shapes satisfying infinitesimal mass transportation at each point. The result is that, not only the spatial map induces a sub-vertex correspondence between the surfaces, but also the transportation of the whole surface, and thus the bijectivity of the induced map is assured. The performance of the proposed method is assessed on several popular benchmarks.

1. Introduction

The problem of non-rigid shape matching, which is defined as the problem of finding a meaningful map between two surfaces, is ubiquitous in computer vision and graphics. A challenging setting for this problem is the one that involves the application of such methods to shapes that undergo non-rigid transformations, such as articulated mod-

els in different poses. The solution to this problem has many applications. For instance, dense maps between surfaces are often used in the definition of retrieval algorithms [10, 11, 9, 5, 24, 2], to perform shape interpolation [14] and to define statistical models of shapes [12, 3].

Several methods have been proposed over the years by the research community to tackle this problem [30]. Shape matching approaches can be split into two main categories: sparse and dense. Methods that belong to the first category aim at finding a low number of matches which favour quality over quantity of matches, while the latter try to find the whole map between the surfaces.

A recent ground-breaking approach to the estimation of dense correspondences is represented by the functional maps framework, proposed by Ovsjanikov *et al.* [20]. In their work, the authors shift from studying the problem in terms of point-wise mappings to a definition of maps between functional spaces. In particular, they define the functional correspondences in terms of a continuous linear operator that maps functions defined on a surface onto functions defined on the target surface. In this setting, the classical point-wise representation can be seen as a special case in which the functional map is used to transfer delta-functions onto delta-functions. The popularity of this approach is due to its efficiency for representing and computing correspondences and its compactness under specific assumptions. One of the drawbacks of this approach is that the correspondences are only intrinsically captured and the induced functional transformation might not be reducible to a continuous point-wise map, much less one that is guaranteed to be bijective. This, together with the low-rank approximation of the function transformations, results in a lack of localization of the mapping, where impulse functions on one surface are mapped onto *diffuse* functions on the target.

Different refinement steps have been proposed to im-

prove ex-post the localization of the correspondences and reconstruct the point-wise map from the functional representation [25, 8]. Finally, the framework has been extended to deal with non-isometric transformations [21, 17, 26, 16], partial matching [23, 18, 7] and many other fields of application [27, 13, 31].

In this paper, we propose a novel approach to the computation of a continuous bijective map between two surfaces. In particular, we move from the low rank spectral representation of the correspondences between surfaces, *i.e.* the *truncated* functional map, to a sparse spatial representation, *i.e.* the *spatial map* between the surfaces. The approach is based on the observation that continuity and smoothness of the optimal map induces structure both on the spectral and the spatial domain, the former providing effective low rank approximations, while the latter exhibiting strong sparsity and locality that can be used in the solution of large-scale problems. The optimization problem is defined in terms of distortion minimization of functions’ transportation between surfaces. More specifically, we adopt a Gromov-Wasserstein type distance between the transported functions, which in the discrete setting results in a spatial mapping satisfying infinitesimal mass transportation at each point. The term acts both as a data-fidelity term and a commutativity regularizer, further improving the quality of the resulting map. The main benefit of the proposed solution is twofold. First, the spatial map induces sub-vertex correspondence map between surfaces. Second, by enforcing bidirectional infinitesimal area-transport, the formulation guarantees continuity and bijectivity of the induced map.

2. Related work

In this paper, we propose a novel approach to the computation of a continuous bijective map between two surfaces. Our solution induces a sub-vertex correspondence and it assures the complete transportation of the whole surface, and thus bijectivity. In this section, we focus on works relevant to our discussion. We refer to [30] for a comprehensive survey of correspondence algorithms.

Most early methods designed to solve the non-rigid shape matching problem focus on the definition of an optimization problem which minimizes a distortion measure between the given surfaces. Kim *et al.* [15] propose an automatic pipeline to find an intrinsic map between two non-isometric surfaces. In their work, they propose to blend together different maps in order to find a solution that exhibits a low distortion on the matches of pairs of surfaces differing by large deformations. The key idea is that it is possible to tailor different low-dimensional maps for different regions of the surfaces, while the global map can be expressed as a weighted combination of them.

In a recent seminal work, Ovsjanikov *et al.* [20] introduced the *functional map framework*. The authors propose

to model correspondences between surfaces intrinsically by modelling the transformation of functions defined over them. In particular, the mapping is reduced to a linear operator which maps functions defined on a surface into functions on the target surface, which is then represented in a low-rank truncation of the functional basis. This approach, and several other approaches based on it (*i.e.* [18, 23, 21]), exploits the fact that it is often easier to obtain corresponding functions rather than corresponding points. One of the major selling points of this work is that, given a low-rank functional basis, the functional representation can be efficiently computed via linear least squares. Its main drawback, on the other hand, is that it only captures correspondences intrinsically, and the estimated functional correspondence might not correspond to a continuous point-to-point map between the surfaces; the underlying point relation might be neither injective nor onto, and it might not represent a point-to-point map at all. Recent research addresses the problem of point-wise map recovery mainly via ex-post refinement. Rodolà *et al.* [25] propose to cast the point-to-point map recovery as a probability density estimation problem to obtain both a better distance measure (w.r.t. the ℓ^2) and a tool to impose regularity assumptions on the alignment map. In the proposed analysis, the authors show that it is possible to lift the limiting requirement of the input shapes being nearly-isometric, which is one of the main assumptions in the original approach. Gasparetto *et al.* [8] propose to enforce bijectivity by adding constraints mapping delta functions onto smoothed versions of the delta functions.

The crucial difference between the proposed approach and the original functional map framework (FM) is the choice of the domain. Specifically, while both FM and our approach formulate correspondence as a linear map between spaces of functions on the shapes, in FM the map is computed in the *spectral* domain (Laplacian eigenbasis), while the proposed approach produces a map directly in the *spatial* domain. The clear advantage is that unlike FM, we do not need an extra step of spatial map recovery, which adds a significant additional cost and constitutes a challenging problem on its own [8, 25].

Nogneng and Ovsjanikov [19] recently proposed a new data term that guides the optimization process towards functional maps that are closer to point-to-point maps. The authors encode function preservation via commutativity with an underlying map, *i.e.* $\|CA - BC\|^2$ (where C is the truncated map, while A and B are the spectral coefficients of surface descriptors). Differently from the term used in the optimization problem proposed in Section 4, here the function preservation is achieved via spectral coefficients preservation. Finally, Solomon *et al.* [29] propose an algorithm to compute probabilistic correspondences optimizing a Gromov-Wasserstein objective. They define the problem




Figure 1. The image shows how the variance changes through different iterations. On the left, the source mesh with a delta function at one point. On the right, the target mesh with the same delta function projected through the spatial map at different iterations, the first being the projection obtained from the spectral functional map.

as an optimal transportation problem, using the Kullback-Leibler as a measure of distortion. An entropic regularizer is used to penalize non-sparse solutions.

Contributions. Our contributions can be summarized as follows:

- We define the correspondence problem as an optimal transport problem with a variance term which assures bijectivity and sparsity of the solution; the data-fidelity term used in the optimization problem is defined on the spatial domain, and assures the commutativity of the solution.
- The resulting spatial map is a fuzzification of an area-preserving diffeomorphism, while the induced sparsity can be exploited in large scale problems. Moreover, as a result of the induced sparsity and locality, the spatial map determines a continuous dense intrinsic *sub-vertex* correspondence between surfaces. To our knowledge, no other method has this property.
- We show how the proposed approach is able to deal with the challenging setting of partial shape matching.

We want to stress the fact that the proposed approach does not rely on a particular initialization. Any dense correspondence map can be used to build the initial smooth permutation matrix which will be optimized by the proposed method.

Finally, in order to assess the accuracy of the retrieved map, we evaluate our method on standard non-rigid shape matching datasets and on partial shape matching datasets, showing that the proposed approach is both flexible and able to achieve good performance even in the very challenging setting of partial matching.

3. Background

We model shapes as compact smooth connected Riemannian manifolds equipped with an *intrinsic metric* d and the standard measure induced by the volume form.

We define a *correspondence* between two manifolds \mathcal{N} and \mathcal{M} as a bijective function between the manifolds. Adding continuity and differentiability to the mix, we obtain a requirement that the correspondence be a diffeomorphism between the manifolds, *i.e.*, a differentiable and invertible function whose inverse is also differentiable. Area preservation further imposes the volume of the differential of the diffeomorphism to be one at each point.

We first relax the binary notion of correspondence into a *fuzzy* notion introducing a *coupling* function $u : \mathcal{N} \times \mathcal{M} \rightarrow [0, 1]$ such that for every measurable subset $A \subseteq \mathcal{N}$ and $B \subseteq \mathcal{M}$,

$$\int_A \int_{\mathcal{M}} u dy dx = \int_A dx; \quad \int_B \int_{\mathcal{N}} u dx dy = \int_B dy. \quad (1)$$

In other words, $u(x, y) dx dy$ defines a weighted product measure on $\mathcal{N} \times \mathcal{M}$ whose marginals are the measures dx and dy on \mathcal{N} and \mathcal{M} , respectively. The quantity $u(x, y) dx$ can be thought of as the infinitesimal amount of mass transported from point x on \mathcal{N} to point y on \mathcal{M} . The mapping can be interpreted as a probabilistic assignment since for every $x \in \mathcal{N}$, $u(x, y)/|\mathcal{M}|$ forms a probability distribution over \mathcal{M} , where $|\mathcal{M}| = \int_{\mathcal{M}} dy$. Conversely, for every $y \in \mathcal{M}$, $u(x, y)/|\mathcal{N}|$ represents a distribution over \mathcal{N} .

Functional correspondence. Given a manifold \mathcal{M} , let $f, g : \mathcal{M} \rightarrow \mathbb{R}$ be real scalar fields on the manifold. We define the standard inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} f(x)g(x)dx$, where integration is done using the manifold's Riemannian metric. Endowed with this scalar product and the resulting

norm, we can define the space of square-integrable functions over \mathcal{M} , denoted as

$$L^2(\mathcal{M}) = \{f : \mathcal{M} \rightarrow \mathbb{R} \mid \langle f, f \rangle_{\mathcal{M}} < \infty\}. \quad (2)$$

Given two manifolds \mathcal{N} and \mathcal{M} , Ovsjanikov *et al.* [20] proposed modelling the correspondences between the spaces of square-integrable functions $L^2(\mathcal{N})$ and $L^2(\mathcal{M})$ in terms of a linear operator $T : L^2(\mathcal{N}) \rightarrow L^2(\mathcal{M})$ mapping functions over \mathcal{N} onto functions over \mathcal{M} . Clearly this setting generalizes classical vertex-wise correspondences as this can be achieved by a T that maps delta-functions onto delta-functions. However, also non-vertex-wise correspondences arise from the formulation.

Given bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ on $L^2(\mathcal{N})$ and $L^2(\mathcal{M})$ respectively, the functional correspondence can be expressed w.r.t. to these bases as follows:

$$\begin{aligned} Tf &= T \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} \phi_i = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} T \phi_i \\ &= \sum_{ij \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} \underbrace{\langle T \phi_i, \psi_j \rangle_{\mathcal{M}}}_{c_{ij}} \psi_j, \end{aligned} \quad (3)$$

where the coefficients c_{ij} depend on the choice of the bases. Taking only k elements of each basis, one obtains a rank- k approximation of T as a $k \times k$ matrix $\mathbf{C} = (c_{ij})$.

In order to compute \mathbf{C} , Ovsjanikov *et al.* [20] assume to be given a set of q corresponding functions $\{f_1, \dots, f_q\} \subseteq L^2(\mathcal{N})$ and $\{g_1, \dots, g_q\} \subseteq L^2(\mathcal{M})$. Denoting by $a_{ij} = \langle f_j, \phi_i \rangle_{\mathcal{N}}$ and $b_{ij} = \langle g_j, \psi_i \rangle_{\mathcal{M}}$ the $k \times q$ matrices of the respective coefficients onto the selected bases, functional correspondence boils down to the linear system

$$\mathbf{C}\mathbf{A} = \mathbf{B}. \quad (4)$$

If $q \geq k$, the system (4) is (over-)determined and is solved in the least squares sense to estimate \mathbf{C} .

Ovsjanikov *et al.* [20] also showed that it is convenient to use the eigenfunctions of the Laplace-Beltrami operators of \mathcal{N} and \mathcal{M} as the bases $\{\phi_i, \psi_i\}_{i \geq 1}$, since under these bases truncating the series at the first k coefficients has the effect of ‘‘low-pass’’ filtering over the functional representations. Further, expressed in the Fourier bases, the matrix \mathbf{C} has interesting properties making it more efficient to estimate.

Discretization. In the discrete setting, the manifold \mathcal{N} is sampled at n points x_1, \dots, x_n which are connected by edges E and faces F , forming a manifold triangular mesh (V, E, F) . Similarly, \mathcal{M} is sampled at m points y_1, \dots, y_m . In this setting, a function on the manifold is represented by an n -dimensional vector $\mathbf{f} = (f(x_1), \dots, f(x_n))^T$. The inner product is discretized as $\langle \mathbf{f}, \mathbf{g} \rangle = \mathbf{f}^T \mathbf{S} \mathbf{g}$, where $\mathbf{S} = \text{diag}(s_1, \dots, s_n)$ and $s_i = \frac{1}{3} \sum_{j,k:ijk \in F} A_{ijk}$ denotes the local area element.

Similarly, the coupling function u is discretized into a $m \times n$ dimensional matrix $\mathbf{U} = (u_{ij})$ with $u_{ij} = u(x_j, y_i)$ for which the following holds:

$$\begin{aligned} \mathbf{s}_{\mathcal{M}}^T \mathbf{U} &= (\mathbf{s}_{\mathcal{N}}^T \mathbf{1}) \mathbf{1}^T = |\mathcal{N}| \mathbf{1}^T \\ \mathbf{U} \mathbf{s}_{\mathcal{N}} &= (\mathbf{1}^T \mathbf{s}_{\mathcal{M}}) \mathbf{1} = |\mathcal{M}| \mathbf{1} \end{aligned} \quad (5)$$

where $\mathbf{s}_{\mathcal{N}}$ and $\mathbf{s}_{\mathcal{M}}$ are the vectors of the the area elements of \mathcal{N} and \mathcal{M} respectively.

Note that, in the discrete domain, in general there can be no bijective map between the vertices on the two meshes. However, \mathbf{U} can fully capture the many-to-many mapping between the area elements associated to each vertex. In fact, let $\theta : \mathcal{N} \rightarrow \mathcal{M}$ be an area-preserving diffeomorphism between \mathcal{N} and \mathcal{M} , the corresponding coupling is simply the integration over the area elements of the operator associated with θ . In other words, the associated continuous and discrete couplings are

$$u(x, y) = \delta_{\theta(x)}(y) \quad (6)$$

$$u_{ij} = \int_{A_i} \int_{A_j} u(x, y) dx dy \quad (7)$$

where A_i is the surface patch around vertex i and $\delta_x(y)$ is the Dirac delta function having zero integral and satisfying $\delta_x(y) = 0$ for $y \neq x$. It is easy to show that the function $u(x, y)$ thus defined is indeed a coupling:

$$\begin{aligned} \int_{\mathcal{M}} \delta_{\theta(x)}(y) dy &= 1 \\ \int_{\mathcal{N}} \delta_{\theta(x)}(y) dx &= \int_{\mathcal{M}} \delta_{\theta^{-1}(y)}(s) \frac{ds}{|D\theta(s)|} = 1 \end{aligned}$$

where $|D\theta(s)|$ is the volume of the differential of the diffeomorphism θ , which is equal to 1 since the diffeomorphism is assumed to be area-preserving.

From the coupling function U , we can define the mapping posteriors over \mathcal{N} and \mathcal{M} respectively as

$$\mathbf{\Pi} = \frac{1}{|\mathcal{N}|} \mathbf{U} \mathbf{s}_{\mathcal{N}} \quad \mathbf{\Omega} = \frac{1}{|\mathcal{M}|} \mathbf{s}_{\mathcal{M}} \mathbf{U}.$$

For the mapping posteriors, the following holds:

$$\mathbf{\Pi} \mathbf{1} = \mathbf{1}, \quad \mathbf{s}_{\mathcal{M}}^T \mathbf{\Pi} = \mathbf{s}_{\mathcal{N}}^T, \quad \mathbf{1}^T \mathbf{\Omega} = \mathbf{1}^T, \quad \mathbf{\Omega} \mathbf{s}_{\mathcal{N}} = \mathbf{s}_{\mathcal{M}}^T \quad (8)$$

4. Problem formulation

Starting from the observation that coupling functions can be used to effectively transport functions from one manifold onto the other, we pose the functional mapping in the spatial domain. In fact, given a coupling u , we can reconstruct the expectation over the transportation of $f : \mathcal{N} \rightarrow \mathbb{R}$ as

$$\hat{g} = \frac{1}{|\mathcal{N}|} \int_{\mathcal{N}} u(x, y) f(x) dx. \quad (9)$$

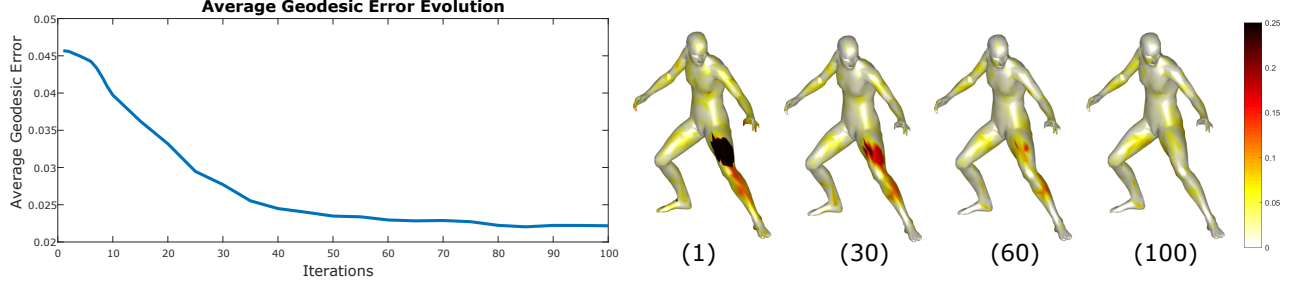


Figure 2. The graph on the left shows the evolution of the average geodesic error computed on the mesh over the first 100 iterations of the proposed method. On the right we show the point-wise geodesic errors computed at iterations 1, 30, 60 and 100.

Conversely, a function $g : \mathcal{M} \rightarrow \mathbb{R}$ can be mapped back onto \mathcal{N} as

$$\hat{f} = \frac{1}{|\mathcal{M}|} \int_{\mathcal{M}} u(x, y) g(y) dy. \quad (10)$$

In the discrete domain we can rewrite these relations as

$$\hat{\mathbf{g}} = \frac{1}{|\mathcal{N}|} \mathbf{U} \mathbf{S}_{\mathcal{N}} \mathbf{f} = \mathbf{\Pi} \mathbf{f}, \quad (11)$$

$$\hat{\mathbf{f}}^T = \mathbf{g}^T \frac{1}{|\mathcal{M}|} \mathbf{S}_{\mathcal{M}} \mathbf{U} = \mathbf{g}^T \mathbf{\Omega}. \quad (12)$$

For the rest of the paper we will assume the shapes to be area-normalized and thus drop the terms $\frac{1}{|\mathcal{N}|}$ and $\frac{1}{|\mathcal{M}|}$.

The spatial functional map can thus be estimated by searching for coupling matrix \mathbf{U} that better transport a set of functions $\{f_i\}_{i=1, \dots, k}$ over \mathcal{N} onto the corresponding functions $\{g_i\}$ over \mathcal{M} . Here we follow Nogneng and Ovsjanikov's commutative descriptor formulation [19] resulting in the data term

$$E_{\text{data}} = \sum_{i=1}^k \|\mathbf{U} \mathbf{S}_{\mathcal{N}} \text{diag}(\mathbf{f}_i) - \text{diag}(\mathbf{g}_i) \mathbf{S}_{\mathcal{M}} \mathbf{U}\|. \quad (13)$$

Using the same commutative functional formulation, we additionally introduce a metric distortion term

$$E_{\text{dist}} = \|\mathbf{U} \mathbf{S}_{\mathcal{N}} \mathbf{D}_{\mathcal{N}} - \mathbf{D}_{\mathcal{M}} \mathbf{S}_{\mathcal{M}} \mathbf{U}\|, \quad (14)$$

where $\mathbf{D}_{\mathcal{N}}$ and $\mathbf{D}_{\mathcal{M}}$ are the distance matrices in \mathcal{N} and \mathcal{M} respectively.

In order to guarantee that the fuzzy coupling is as close as possible to an underlying area-preserving diffeomorphism, we also require the couplings to have mapping posteriors that are as tightly concentrated around as few spatially-coherent vertices as possible. We do this by adding penalty terms related to the manifold variance of both the forward and backward mapping posteriors $\mathbf{\Pi}$ and $\mathbf{\Omega}$. Let $p : \mathcal{M} \rightarrow \mathbb{R}$ be a distribution over a manifold \mathcal{M} , the Karcher mean of p is the manifold point μ such that

$$\mu = \underset{y}{\operatorname{argmin}} \int_{\mathcal{M}} d^2(x, y) p(x) dx \quad (15)$$

with the variance of the distribution corresponding to the attained minimum

$$\operatorname{Var}(p) = \min_y \int_{\mathcal{M}} d^2(x, y) p(x) dx. \quad (16)$$

Here we propose a quadratic approximation of the variance:

$$V(p) = \int_{\mathcal{M}} \int_{\mathcal{M}} d^2(x, y) p(x) p(y) dx dy. \quad (17)$$

Theorem 1. *Let \mathcal{M} be a Riemannian manifold endowed with a measure, and let $p : \mathcal{M} \rightarrow \mathbb{R}$ be a probability density distribution over \mathcal{M} , the following holds:*

1. $\operatorname{Var}(p) \leq V(p) \leq 4 \operatorname{Var}(p)$
2. $\lim_{\operatorname{Var}(p) \rightarrow 0} \frac{V(p)}{\operatorname{Var}(p)} = 2$

Proof. The proof can be found in the supplementary material. \square

Thanks to this result we can use $V(p)$ as a valid proxy of $2 \operatorname{Var}(p)$, especially since we are interested in low variance solutions. This leads to the following variance terms both in \mathcal{N} and \mathcal{M} :

$$V_{\mathcal{M}}(u) = \int_{\mathcal{N}} \int_{\mathcal{M}} \int_{\mathcal{M}} d_{\mathcal{M}}^2(x, y) u(z, x) u(z, y) dx dy dz$$

$$V_{\mathcal{N}}(u) = \int_{\mathcal{M}} \int_{\mathcal{N}} \int_{\mathcal{N}} d_{\mathcal{N}}^2(x, y) u(x, z) u(y, z) dx dy dz$$

These penalty terms are easily discretized as the quadratic forms

$$V_{\mathcal{M}} = \operatorname{Tr}(\mathbf{U}^T \mathbf{S}_{\mathcal{M}} \mathbf{D}_{\mathcal{M}}^2 \mathbf{S}_{\mathcal{M}} \mathbf{U} \mathbf{S}_{\mathcal{N}}) \quad (18)$$

$$V_{\mathcal{N}} = \operatorname{Tr}(\mathbf{U} \mathbf{S}_{\mathcal{N}} \mathbf{D}_{\mathcal{N}}^2 \mathbf{S}_{\mathcal{N}} \mathbf{U}^T \mathbf{S}_{\mathcal{M}}). \quad (19)$$

In a formulation strongly related to ours, Solomon *et al.* [29] adopt a similar spatial representation where the couplings are sought to minimize the metric distortion in terms

of the Gromov-Wasserstein distance between the Riemannian Manifolds. Here too the authors are interested in extracting couplings that are as deterministic as possible and for this reason they introduce an entropic term that penalizes posterior mappings with large support. While this has similar sparsifying effects as our variance-based regularizer, it has no knowledge of geometric proximity, which is central to the definition of variance. Hence, it does *not* impose locality and compactness of the support in the same way that our terms do. Further, by imposing bidirectional constraints, we enforce the representation to be closer to a many-to-many map actually induced by a continuous bijective mapping between the underlying manifolds.

Putting the terms together, we obtain the following formulation:

$$\begin{aligned} \underset{\mathbf{U}}{\text{minimize}} \quad & \sum_{i=1}^k \|\mathbf{U}\mathbf{S}_{\mathcal{N}} \text{diag}(\mathbf{f}_i) - \text{diag}(\mathbf{g}_i)\mathbf{S}_{\mathcal{M}}\mathbf{U}\| \quad (20) \\ & + \alpha \|\mathbf{U}\mathbf{S}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}} - \mathbf{D}_{\mathcal{M}}\mathbf{S}_{\mathcal{M}}\mathbf{U}\| \\ & + \beta_1 \text{Tr}(\mathbf{U}^T\mathbf{S}_{\mathcal{M}}\mathbf{D}_{\mathcal{M}}^2\mathbf{S}_{\mathcal{M}}\mathbf{U}\mathbf{S}_{\mathcal{N}}) \\ & + \beta_2 \text{Tr}(\mathbf{U}\mathbf{S}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}}^2\mathbf{S}_{\mathcal{N}}\mathbf{U}^T\mathbf{S}_{\mathcal{M}}) \end{aligned}$$

$$\begin{aligned} \text{subject to} \quad & \mathbf{U} \geq 0 \\ & \mathbf{U}\mathbf{s}_n = \mathbf{1} \\ & \mathbf{U}^T\mathbf{s}_m = \mathbf{1} \end{aligned}$$

where α , β_1 , and β_2 are coefficients used to balance the effects of the penalty terms.

4.1. Dealing with partiality

In order to allow partiality, and, in particular, to be able to map partial shapes to a reference full shape as in [23], we relax the bidirectional area transportation constraint from the coupling, allowing for all $y \in \mathcal{M}$

$$0 \leq \int_{\mathcal{N}} u(x, y) dx \leq 1. \quad (21)$$

The upper bound is further relaxed by adding a hinge loss penalty term

$$\int_{\mathcal{M}} \left(\int_{\mathcal{N}} u(x, y) dx - 1 \right)_+ dy, \quad (22)$$

where $(x)_+ = \max(0, x)$.

This results in the reduced set of constraints in the discretized problem

$$\begin{aligned} \mathbf{U} & \geq 0 \\ \mathbf{U}^T\mathbf{s}_m & = \mathbf{1} \end{aligned}$$

which, by substituting for the mapping posterior $\Omega = \mathbf{S}_{\mathcal{M}}\mathbf{U}$ reduces to the standard multi-simplex

$$\begin{aligned} \Omega & \geq 0 \\ \Omega^T\mathbf{1} & = \mathbf{1}. \end{aligned}$$

Under this re-writing, the formulation becomes

$$\begin{aligned} \underset{\Omega}{\text{minimize}} \quad & \sum_{i=1}^k \|\mathbf{S}_{\mathcal{M}}^{-1}\Omega\mathbf{S}_{\mathcal{N}} \text{diag}(\mathbf{f}_i) - \text{diag}(\mathbf{g}_i)\Omega\| \quad (23) \\ & + \alpha \|\mathbf{S}_{\mathcal{M}}^{-1}\Omega\mathbf{S}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}} - \mathbf{D}_{\mathcal{M}}\Omega\| \\ & + \beta_1 \text{Tr}(\Omega^T\mathbf{D}_{\mathcal{M}}^2\Omega\mathbf{S}_{\mathcal{N}}) \\ & + \beta_2 \text{Tr}(\mathbf{S}_{\mathcal{M}}^{-1}\Omega\mathbf{S}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}}^2\mathbf{S}_{\mathcal{N}}\Omega^T) \\ & + \gamma \mathbf{s}_{\mathcal{M}}^T (\mathbf{S}_{\mathcal{M}}^{-1}\Omega\mathbf{S}_{\mathcal{N}} - \mathbf{1})_+ \\ \text{subject to} \quad & \Omega \geq 0 \\ & \Omega^T\mathbf{1} = \mathbf{1}. \end{aligned}$$

4.2. Scaling behaviour and sparsity

One problem shared by most spatial approaches is their scaling behaviour: having to optimize the coupling, the space requirement grows quadratically with the number of vertices, while even just computing the objective function can scale cubically. This clearly limits the resolution and complexity of the meshes that can be analysed. The spectral representation of the functional map addresses this limitation by adopting a low-rank approximation in the Fourier basis, which allows us to tame both spatial and time complexity for extracting the solution.

In our approach we deal with the scaling behaviour by making use of the sparsity of the coupling representation in the spatial domain when in the vicinity of the solution. Assuming the existence of an underlying area-preserving diffeomorphism between the two manifolds, equation (7) tells us that the coupling matrix \mathbf{U} associated with the optimal diffeomorphism must not only be sparse, but the support of each row and column must be compact and local in \mathcal{N} and \mathcal{M} respectively.

Further, when sufficiently close to the optimal coupling, the variance terms induce more compact and local support. This results in the fact that if we start from a coupling that is both compact and sufficiently close to the optimal solution, not only would the sparsity and compactness increase, but we can safely expect no new elements to enter the support. This allows us to only work on a low-dimensional face of the multi-simplex, thus reducing the space and time complexity of the approach.

To this end, we start from a solution \mathbf{C} obtained with a spectral functional map approach. From this we can obtain the spatial mapping $\mathbf{U} = \Psi\mathbf{C}\Phi^T$ which can then be thresh-

olded to select which couplings to force to zero in the optimization process to sparsify the representation. Finally, an initialization point for the optimization can be obtained by projecting the sparsified \mathbf{U} onto the discrete coupling constraints. Here either least square or Sinkhorn [28] projection can be used.

Once we have induced a sparse representation for \mathbf{U} , the whole optimization process scales down. Indeed, the area element matrices $\mathbf{S}_{\mathcal{N}}$ and $\mathbf{S}_{\mathcal{M}}$ being diagonal are clearly sparse and the multiplication by them scales linearly with the number of couplings. A more interesting analysis comes from the geodesic distance matrices $\mathbf{D}_{\mathcal{N}}$ and $\mathbf{D}_{\mathcal{M}}$. Here the compactness of the solution plays an essential role in the ability to sparsify the computation of the variance terms: since the marginal mappings obtained from rows and columns of \mathcal{U} have local, compact support, only the distances within the radius of such support enter in the computation of the variances. Thus not all the entries of the geodesic distance matrices need to be computed, but rather we can maintain a local sparse representation easily computed with a marching front approach.

The metric distortion term $\|\mathbf{U}\mathbf{S}_{\mathcal{N}}\mathbf{D}_{\mathcal{N}} - \mathbf{D}_{\mathcal{M}}\mathbf{S}_{\mathcal{M}}\mathbf{U}\|$ on the other hand, has neither local nor sparse support and quickly becomes a bottleneck in the scaling behaviour of the approach. We can solve this by substituting the distance operators with any other metric-derived operator with compact, local support; for example a thresholded Gaussian geodesic kernel $K = (k_{ij})$ with

$$k_{ij} = \begin{cases} \exp(-\frac{1}{2} \frac{d_{ij}^2}{\sigma^2}) & \text{if } d_{ij} < R \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

where d_{ij} is the geodesic distance between vertices i and j , and R is a the maximum radius of the marginal supports.

5. Experimental evaluation

We compare the spatial map computed as the result of our optimization problem with the current state-of-the-art approaches dealing with the non-rigid shape matching problem. The first method we compare against is the functional map framework [20] (*FM* in the graphs' legend), whose results are computed using the iterative refinement procedure proposed in the original work. Then, we compare with the two refinement processes [25, 8] (*CPD* and *InjCorr* respectively) introduced in Section 2. Finally, we compare with the work of Nogneng and Ovsjanikov [19] (*IDPC*). They recently proposed a new approach to the computation of the functional map between two surfaces using a commutative descriptor preservation term.

We evaluate the methods quantitatively and qualitatively on the TOSCA [1] and SCAPE [3] datasets. The former contains 76 meshes distributed in 9 classes, both humanoid



Figure 3. Qualitative results on the challenging setting of partial matching. The null mesh is in the first column, while the second shows the dense correspondence induced by the *Partial Functional Map* work. The third column shows the refined map produced by the proposed approach.

and non-humans. The latter contains 71 meshes representing the same individual in different poses, and is the result of the application of a full-body scanning procedure. The meshes representing the same class in both dataset are in point-wise correspondence. For this reason, and for efficiency purposes, we sub-sampled *non-coherently* all the meshes. As a result, the processed meshes are no more in point-to-point correspondence (and differ in the number of vertices). Finally, using the full-resolution meshes as proxies, we compute the barycentric coordinates of each point with respect to the other meshes, thus leading to a *sub-vertex* ground-truth correspondence map between the surfaces. Thanks to this asynchronous simplification process, we were able to create a high-quality dataset with each mesh having a small number of vertices (≈ 3000). This dataset will be made publicly available, along with the implementation of our approach. Furthermore, the methods have been tested also on a dataset in which the vertex-to-vertex correspondence among the shapes have been preserved. Due to lack of space, the performance achieved by the methods has been included in the supplementary material.

As a measure of error, we compute the geodesic distance between the ground-truth and the correspondence map produced by the different methods. The proposed approach yields the spatial map Ω as a result, in which each column represents a *fuzzy assignment*. Sub-vertex correspondences are extracted by computing the Karcher mean over the soft-assignments.

Note that, since the variance term in the proposed optimization process induces sparsity on the computed spatial map, the resulting soft-assignment concentrates the transported mass around a small number of points on the target surface. In the common scenario in which most of the mass is concentrated on the points of a single triangle of the

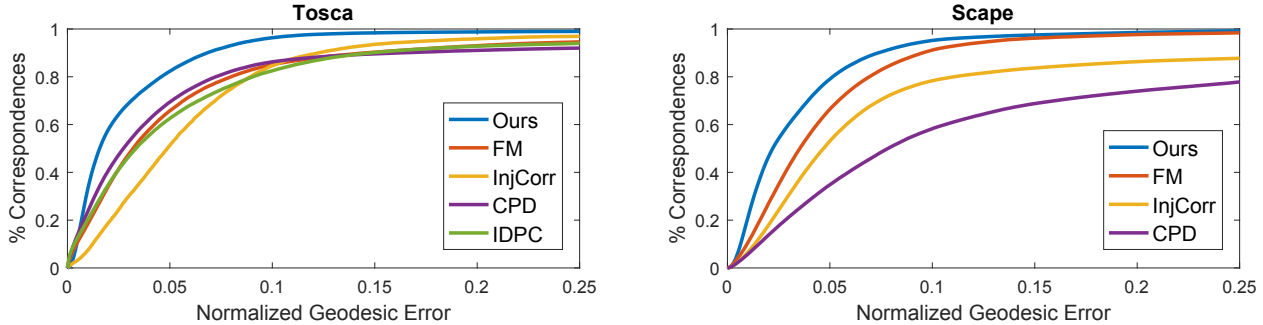


Figure 4. Comparison of the normalized error curves obtained on the re-sampled SCAPE and TOSCA datasets using several shape matching algorithms, while allowing the symmetric solutions.

mesh, the Karcher mean can be approximated as weighted barycentric coordinate, where the weights are the normalized probabilities associated to the points of the triangle.

Comparisons. The functional maps used in the comparisons both on the original pipeline and on the refinement approaches, are constructed by solving a least-squares system $\mathbf{CA} = \mathbf{B}$, where \mathbf{A} and \mathbf{B} are the spectral coefficients. In the original work, Ovsjanikov *et al.* [20] show that adding descriptors which are in correspondence greatly improves the quality of the functional map. In this comparison, the initial correspondence is established using a sparse matching algorithm [22] (≈ 200 matches between shapes). The functional map is computed with $k = 75$ basis functions. The wave kernel signature [4] is used both in the function preservation term introduced in Section 4, as well as for the computation of the sparse matching via [22].

Figure 2 shows the mean geodesic error as it changes over several iterations of the proposed method. In particular, the curve on the graph shows the average geodesic error evolution, while the meshes are rendered using the point-wise geodesic error as the color-map. Note that the geodesic error on the fingers of the humanoid (which are in general hard to match correctly, in particular with methods based on intrinsic measures) has been corrected in the first iterations. Similarly, Figure 1 shows the evolution of the fuzzy mapping of a vertex onto the target shape. It is quite evident how effective our approach is at improving the localization of the map. In Figure 4, we show quantitative results yielded by the compared methods on both TOSCA and SCAPE datasets. In particular, the graphs show the geodesic error curves, which represent the percentage of points whose normalized geodesic error is below a certain threshold. From the plotted curves we can see that the proposed method generally outperforms the compared ones.

Partial matching. To assess the performance on the partial matching setting, we apply our method to the SHREC’16 *holes* dataset [6]. The shapes included in the dataset span different classes and are based on the TOSCA high-resolution dataset. The dataset includes ‘null’ shapes for

each class, *i.e.* a full shape in a canonical pose. Following the adaptation of the proposed method introduced in Section 4, we perform a part-to-whole matching between a partial shape and its full counterpart. The dataset is provided with ground-truth vertex-to-vertex correspondence between partial and full meshes. Since dealing with the partial setting is not the focus of this paper, we present only some qualitative results. Figure 3 shows these results. Our method was initialized with a functional map computed using [23]. In particular, the image shows in the first column the full mesh, while in the second we show the dense correspondence induced by the partial functional maps pipeline [23]. The third column presents the correspondence map obtained with the proposed approach.

6. Conclusion

In this paper we proposed a novel solution to the non-rigid shape matching problem. We formulated the problem in terms of a functional correspondence in the spatial domain, by optimizing how a fuzzy mapping between the vertices transforms functions defined over one shape onto corresponding functions over the other shape. One fundamental characteristic of the approach is that it does not assume the existence of vertex-to-vertex correspondences, but rather can extract a *sub-vertex* mapping onto the target shape. We introduced a novel variance-based regularizer that enforces both sparsity and locality of the marginal vertex-correspondences, which, by adopting the loosely localized initialization provided by spectral functional maps, allows us to limit the scaling behaviour inherent in spatial approaches by making use of the intrinsic sparsity and locality of the support of the coupling function. Further, by relaxing the coupling constraints in one direction, the approach was shown to be able to cope with partiality. Extensive comparative evaluations show that our approach provides better-localized mapping competitive approaches at the state of the art both on datasets without vertex-to-vertex correspondences and on datasets with vertex-to-vertex correspondences.

References

- [1] R. K. A. M. Bronstein, M. M. Bronstein. *Numerical geometry of non-rigid shapes*. Springer, 2008. 7
- [2] A. Albarelli, F. Bergamasco, L. Rossi, S. Vascon, and A. Torsello. A stable graph-based representation for object recognition through high-order matching. pages 3341–3344, 2012. cited By 3. 1
- [3] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. *ACM Trans. Graph.*, 24:408–416, 2005. 1, 7
- [4] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops*, pages 1626–1633, Nov 2011. 8
- [5] A. C. Berg, T. L. Berg, and J. Malik. Shape matching and object recognition using low distortion correspondences. In *Proc. CVPR*, 2005. 1
- [6] L. Cosmo, E. Rodolà, M. M. Bronstein, A. Torsello, D. Cremers, and Y. Sahillioğlu. SHREC’16: Partial matching of deformable shapes. In *Proc. 3DOR*, pages 61–67, 2016. 8
- [7] L. Cosmo, E. Rodolà, J. Masci, A. Torsello, and M. M. Bronstein. Matching deformable objects in clutter. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 1–10, Oct 2016. 2
- [8] A. Gasparetto, L. Cosmo, A. Torsello, and R. Wilson. Non-rigid dense bijective maps. In *Proc. ICPR*, 2016. 2, 7
- [9] A. Gasparetto, G. Minello, and A. Torsello. A non-parametric spectral model for graph classification. In *Proceedings of the International Conference on Pattern Recognition Applications and Methods - Volume 1, ICPRAM 2015*, pages 312–319, Portugal, 2015. SCITEPRESS - Science and Technology Publications, Lda. 1
- [10] A. Gasparetto, G. Minello, and A. Torsello. Non-parametric spectral model for shape retrieval. In *Proc. 3DV*, pages 344–352, 2015. 1
- [11] A. Gasparetto and A. Torsello. A statistical model of riemannian metric variation for deformable shape analysis. In *Proc. CVPR*, pages 1219–1228, 2015. 1
- [12] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel. A statistical model of human pose and body shape. *Computer Graphics Forum*, 28(2):337–346, 2009. 1
- [13] Q. Huang, F. Wang, and L. Guibas. Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.*, 33(4):36:1–36:11, July 2014. 2
- [14] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. *ACM Trans. Graph.*, 26(3), July 2007. 1
- [15] V. G. Kim, Y. Lipman, and T. A. Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011. 2
- [16] A. Kovnatsky, M. M. Bronstein, X. Bresson, and P. Vandergheynst. Functional correspondence by matrix completion. *CoRR*, abs/1412.8070, 2014. 2
- [17] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. *CoRR*, abs/1210.0026, 2012. 2
- [18] O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. Non-rigid puzzles. *Computer Graphics Forum*, 35(5):135–143, 2016. 2
- [19] D. Nogneng and M. Ovsjanikov. Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum*, 36(2), 2017. 2, 5, 7
- [20] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4):30:1–30:11, July 2012. 1, 2, 4, 7, 8
- [21] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32(2pt4):459–468, 2013. 2
- [22] E. Rodolà, A. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, pages 182–189, June 2012. 8
- [23] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Computer Graphics Forum*, 36(1):222–236, 2017. 2, 6, 8
- [24] E. Rodolà, L. Cosmo, O. Litany, M. M. Bronstein, A. M. Bronstein, N. Audebert, A. B. Hamza, A. Boulch, U. Castellani, M. N. Do, A.-D. Duong, T. Furuya, A. Gasparetto, Y. Hong, J. Kim, B. L. Saux, R. Litman, M. Masoumi, G. Minello, H.-D. Nguyen, V.-T. Nguyen, R. Ohbuchi, V.-K. Pham, T. V. Phan, M. Rezaei, A. Torsello, M.-T. Tran, Q.-T. Tran, B. Truong, L. Wan, and C. Zou. Deformable Shape Retrieval with Missing Parts. In I. Pratikakis, F. Dupont, and M. Ovsjanikov, editors, *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2017. 1
- [25] E. Rodolà, M. Möller, and D. Cremers. Point-wise map recovery and refinement from functional correspondence. In *Proc. VMV*, 2015. 2, 7
- [26] E. Rodolà, S. Rota Bulò, T. Windheuser, M. Vestner, and D. Cremers. Dense non-rigid shape correspondence using random forests. In *Proc. CVPR*, pages 4177–4184, June 2014. 2
- [27] R. M. Rustamov, M. Ovsjanikov, O. Azencot, M. Ben-Chen, F. Chazal, and L. Guibas. Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph.*, 32(4):72:1–72:12, July 2013. 2
- [28] R. Sinkhorn. A Relationship Between Arbitrary Positive Matrices and Doubly Stochastic Matrices. *The Annals of Mathematical Statistics*, 35(2):876–879, 1964. 7
- [29] J. Solomon, G. Peyré, V. G. Kim, and S. Sra. Entropic metric alignment for correspondence problems. *ACM Trans. Graph.*, 35(4):72:1–72:13, July 2016. 2, 5
- [30] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011. 1, 2
- [31] F. Wang, Q. Huang, and L. J. Guibas. Image co-segmentation via consistent functional maps. In *Proc. ICCV*, pages 849–856, Dec 2013. 2