

Non-Rigid Puzzles

O. Litany^{1,5}, E. Rodolà², A. M. Bronstein^{3,4}, M. M. Bronstein^{2,4}, D. Cremers⁵

¹Tel Aviv University, Israel ²University of Lugano, Switzerland ³Technion, Israel ⁴Intel Perceptual Computing, Israel ⁵TU Munich, Germany




Figure 1: Example of the non-rigid puzzle problem considered in this paper: given a model human shape (leftmost, first column) and three query shapes (two deformed parts of the human and one unrelated ‘extra’ shape of a cat head), the goal is to find a segmentation of the model shape (second column, shown in yellow and green; white encodes parts without correspondence) into parts corresponding to (subsets of) the query shapes. Third column shows the computed correspondence between the parts (corresponding points are encoded in similar color).

Abstract

Shape correspondence is a fundamental problem in computer graphics and vision, with applications in various problems including animation, texture mapping, robotic vision, medical imaging, archaeology and many more. In settings where the shapes are allowed to undergo non-rigid deformations and only partial views are available, the problem becomes very challenging. To this end, we present a non-rigid multi-part shape matching algorithm. We assume to be given a reference shape and its multiple parts undergoing a non-rigid deformation. Each of these query parts can be additionally contaminated by clutter, may overlap with other parts, and there might be missing parts or redundant ones. Our method simultaneously solves for the segmentation of the reference model, and for a dense correspondence to (subsets of) the parts. Experimental results on synthetic as well as real scans demonstrate the effectiveness of our method in dealing with this challenging matching scenario.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Shape Analysis

1. Introduction

Finding correspondence between deformable shapes is one of the cornerstone problems in computer vision and graphics. The ability to establish correspondence between 3D geometric data is a crucial ingredient in a broad spectrum of applications ranging from animation, texture mapping, and robotic vision, to medical imaging and archaeology [vKZHC01]. The deformable shape correspondence problem comes in a variety of flavors and settings. It is common to distinguish between *rigid* and *non-rigid* correspondence depending on whether the shapes are allowed to undergo deformations (in this case, one can further distinguish between isometric or inelastic deformations, or more general non-isometric deformations that

can also change the shape topology). Second, one distinguishes between *full* and *partial* correspondence (in the latter case, one allows for some parts of the shapes to be missing; this setting arises in numerous applications that involve real data acquisition by 3D sensors, inevitably leading to missing parts due to occlusions or partial view). Finally, there is also the difference between *pairwise* and *multiple* correspondence (in the latter case, one tries to establish correspondence between a collection of shapes).

1.1. Related work

Albeit one of the most broadly studied problems in the domain of geometry processing, correspondence is far from being solved, es-

pecially in some challenging settings. We refer the reader to recent survey papers [vKZHC011, BCB15] for an up-to-date review of existing methods.

Rigid partial correspondence problems arising, e.g., in the fusion or completion of multiple 3D scans have been tackled by ICP-like approaches [AMC08, ART15]. Bronstein *et al.* [BB08b] used a regularized ICP approach where the matching parts are explicitly modeled, and proposed a functional similar to the Mumford-Shah [MS89, VC02] imposing part regularity. Litany *et al.* [LBB12] extended this approach to multiple rigid shape matching.

Non-rigid partial correspondence. Several approaches for intrinsic partial matching revolve around the notion of minimum distortion correspondence [BBK06]. Bronstein *et al.* [BB08a, BBBK09] combined metric distortion minimization with optimization over regular matching parts. Rodolà *et al.* [RBA*12, RTH*13] relaxed the regularity requirement by allowing sparse correspondences. Windheuser *et al.* [WSSC11] proposed an integer linear programming solution for dense elastic matching. Sahillioğlu and Yemez [SY14b] proposed a voting-based formulation to match shape extremities, which are assumed to be preserved by the partiality transformation. The aforementioned methods are based on intrinsic metric preservation and on the definition of spectral features, hence their accuracy suffers at high levels of partiality – where the computation of these quantities becomes unreliable due to boundary effects and meshing artifacts.

More recent approaches include the alignment of tangent spaces [BW*14] and the design of robust descriptors for partial matching [vKZH13]. Several works tried to employ machine learning methods to deal with partial matches. Masci *et al.* [MBBV15] introduced Geodesic CNN, a deep learning framework for computing dense correspondences between deformable shapes, providing a generalization of the convolutional networks (CNN) to non-Euclidean manifolds. Wei *et al.* [WHC*15] focused on matching *human* shapes undergoing changes in pose by means of classical CNNs, also tackling partiality transformations.

Dynamic fusion is a particular setting of the problem, referring to non-rigid tracking of depth images produced by 3D sensors. Attempts to extend ICP-based methods to such a setting [LSP08] had limited success due to sensitivity to initialization and to the underlying assumption of small deformations. Recent works [NFS15, DTF*15] generalizing the Kinect fusion approach [NIH*11], were based on volumetric representation of 3D data.

Most of the aforementioned correspondence methods are *point-wise*, i.e., one seeks a mapping between vertices of the underlying shapes. Ovsjanikov *et al.* [OBCS*12] introduced *functional maps*, representing correspondences between functional spaces on the respective shapes. While not intended for partial correspondence, follow-up works [KBBV14] showed that functional maps and similar constructions can handle certain settings with missing parts. Rodolà *et al.* [RCB*16] introduced *partial functional correspondence*, an extension of [OBCS*12] where matched parts are explicitly modeled and regularized in a manner similar to [BB08a, BBBK09]. This method has achieved state-of-the-art performance on the recent SHREC'16 Partial Correspondence benchmark [CRB*16].

Multiple shape correspondence in the rigid settings has been addressed in numerous works, including [HFG*06, TRA11, LBB12]. In the non-rigid setting, pointwise and functional maps for large shape collections have been explored in [HG13, HWG14, SY14a, CRA*16].

1.2. Main contributions

In this paper, we are interested in intrinsic, non-rigid, partial, multiple shape correspondence in a setting which we refer to as *non-rigid puzzles* (see Figure 1). The motivation in mind is to use this formulation as a first step toward automatic reconstruction of the deformable shapes (dynamic fusion), in which one tries to match multiple scans to a near-isometric general model. We assume to be given a model shape and multiple query shapes, assumed to be parts of (near isometrically) deformed versions of the model shape, possibly with additional clutter. The query shapes may contain overlapping parts, and the model shape might have ‘missing’ regions that do not correspond to any query shape; conversely, there might be ‘extra’ query shapes that have no correspondence to the model shape.

We present a framework for solving 3D non-rigid puzzle problems. We formulate such problems as partial functional correspondences between the query and model shapes, and alternate between optimization on the part-to-whole correspondence and the segmentation of the model. Our method can be considered an extension of [RCB*16] for the multiple part setting on one hand, and a non-rigid generalization of the rigid puzzles problem treated in [LBB12] on the other.

The rest of the paper is organized as follows. In Section 2 we overview the basic notions in spectral analysis on manifolds and present the partial functional maps framework. Section 3 formulates the non-rigid puzzle problem and describes the proposed approach, and Section 4 gives the implementation details. Section 5 presents experimental results, where we show how the method copes with some challenging examples. Finally, Section 6 concludes the paper.

2. Background

We model a shape as a two-dimensional Riemannian manifold \mathcal{M} (possibly with boundary $\partial\mathcal{M}$), endowed with the standard measure induced by the volume form. We denote the space of square-integrable functions on the manifold \mathcal{M} by $L^2(\mathcal{M}) = \{f : \mathcal{M} \rightarrow \mathbb{R} \mid \int_{\mathcal{M}} f^2 da < \infty\}$, and use the standard $L^2(\mathcal{M})$ inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} fg da$.

Our manifolds are equipped with the intrinsic gradient $\nabla_{\mathcal{M}}$ and Laplace-Beltrami operator $\Delta_{\mathcal{M}}$, generalizing the corresponding notions from Euclidean spaces to manifolds. By analogy to flat spaces, the Laplacian provides us with all necessary tools for extending Fourier analysis to manifolds. In particular, it admits an eigen-decomposition

$$\Delta_{\mathcal{M}}\phi_i(x) = \lambda_i\phi_i(x) \quad x \in \text{int}(\mathcal{M}) \quad (1)$$

$$\langle \nabla_{\mathcal{M}}\phi_i(x), \hat{n}(x) \rangle = 0 \quad x \in \partial\mathcal{M}, \quad (2)$$

with homogeneous Neumann boundary conditions (2), where \hat{n} is

the normal vector to the boundary. Here, $0 = \lambda_1 \leq \lambda_2 \leq \dots$ are eigenvalues and ϕ_1, ϕ_2, \dots are the corresponding eigenfunctions forming an orthonormal basis of $L^2(\mathcal{M})$.

Since the eigenfunctions of the Laplacian form a basis, any function $f \in L^2(\mathcal{M})$ can be represented via the (manifold) Fourier series expansion

$$f(x) = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i(x). \quad (3)$$

Functional correspondence. A recent paradigm shift in the shape matching problem was introduced by Ovsjanikov *et al.* [OB^{CS}*12]. The authors proposed to model correspondences among two shapes by means of a linear operator $T : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{N})$, mapping functions on \mathcal{M} to functions on \mathcal{N} . Classical point-to-point matching can then be seen as a special case where one maps delta functions to delta functions.

Because T is a linear operator, it can be equivalently represented by a matrix of coefficients $\mathbf{C} = (c_{ij})$ arising from the following short computation: Let us be given orthonormal bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$ on $L^2(\mathcal{M})$ and $L^2(\mathcal{N})$, respectively, and let us fix some function $f \in L^2(\mathcal{M})$. Then

$$\begin{aligned} Tf &= T \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} T \phi_i \\ &= \sum_{i,j \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \underbrace{\langle T \phi_i, \psi_j \rangle_{\mathcal{N}}}_{c_{ij}} \psi_j. \end{aligned} \quad (4)$$

The application of T is expressed by linearly transforming the expansion coefficients of f from basis $\{\phi_i\}_{i \geq 1}$ onto basis $\{\psi_j\}_{j \geq 1}$.

Choosing as the bases the eigenfunctions $\{\phi_i\}_{i \geq 1}$, $\{\psi_j\}_{j \geq 1}$ of the respective Laplacians on the two shapes yields a particularly convenient representation for the functional map [OB^{CS}*12]. By analogy with Fourier analysis, this choice allows to truncate the series (4) after the first k coefficients, which is equivalent to taking the upper left $k \times k$ submatrix of \mathbf{C} as an approximation of the full map. Further, one obtains $c_{ij} = \langle T \phi_i, \psi_j \rangle_{\mathcal{N}} \approx \pm \delta_{ij}$ whenever the two shapes are nearly isometric. This results in matrix \mathbf{C} being diagonally dominant, since $c_{ij} \approx 0$ if $i \neq j$. This particular structure was exploited in [PB^B*13, KB^B*13] as a prior for shape matching problems.

Partial functional correspondence. Let \mathcal{N} be *part* of a shape that is nearly isometric to a full model \mathcal{M} . Recently, Rodolà *et al.* [RCB*16] showed that for each eigenfunction ψ_j of \mathcal{N} there exists a corresponding eigenfunction ϕ_i of \mathcal{M} for some $i \geq j$, such that $c_{ij} = \langle T \phi_i, \psi_j \rangle_{\mathcal{N}} \approx \pm 1$, and zero otherwise. Differently from the full-to-full case of [OB^{CS}*12], where approximate equality holds for $i = j$, here the inequality $i \geq j$ induces a *slanted*-diagonal structure on matrix \mathbf{C} (Figure 2). In particular, the angle of the diagonal can be precomputed and used as a prior for the matching process.

The key idea behind their analysis is to model partiality as a perturbation of the Laplacian matrices $\mathbf{L}_{\mathcal{M}}$, $\mathbf{L}_{\mathcal{N}}$ of the two shapes. Specifically, consider the dog shape \mathcal{M} shown in Figure 3, and assume a vertex ordering where the points contained in the red region




Figure 2: A key observation behind partial functional maps is the eigenfunction interleaving property, by which the eigenvectors of the full shape contain a subset whose restriction constitutes also the (approximate) eigenvectors of a part of the shape. As a result, the inner products of the eigenvectors of the full shape and the part, restricted to the corresponding subset, form a slanted diagonal matrix. The slope of the slant depends on the ratio of the areas of the partial and the full shape.




Figure 3: Structure of Laplacian of a shape consisting of two parts (reproduced from [RCB*16]).

\mathcal{N} appear before those of the blue region $\bar{\mathcal{N}}$. Then, the Laplacian of the full shape $\mathbf{L}_{\mathcal{M}}$ will assume the structure

$$\mathbf{L}_{\mathcal{M}} = \begin{pmatrix} \mathbf{L}_{\mathcal{N}} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_{\bar{\mathcal{N}}} \end{pmatrix} + \begin{pmatrix} \mathbf{P}_{\mathcal{N}} & \mathbf{E} \\ \mathbf{E}^T & \mathbf{P}_{\bar{\mathcal{N}}} \end{pmatrix}, \quad (5)$$

where the second matrix encodes the perturbation due to the boundary interaction between the two regions. Such a matrix is typically very sparse and low-rank, since it contains non-zero elements only in correspondence of the edges connecting $\partial\mathcal{N}$ to $\partial\bar{\mathcal{N}}$.

If the perturbation matrix is identically zero, then (5) is exactly block-diagonal; this describes the case in which \mathcal{N} and $\bar{\mathcal{N}}$ are disjoint parts, and the eigenpairs of $\mathbf{L}_{\mathcal{M}}$ are an interleaved sequence of those of the two blocks. The key result shown in [RCB*16] is that this interleaving property still holds even when considering the full matrix $\mathbf{L}_{\mathcal{M}}$ as given in (5): Its eigenpairs consist of those of the blocks $\mathbf{L}_{\mathcal{N}}$, $\mathbf{L}_{\bar{\mathcal{N}}}$, up to some bounded perturbation that depends on the length and position of the boundary $\partial\mathcal{N}$. In other words, the eigenfunctions and eigenvalues of the parts show up among those

of the full shape. By letting $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$ denote the eigenfunctions on \mathcal{M} and \mathcal{N} respectively, this is what makes the equality $c_{ij} = \langle T\phi_i, \psi_j \rangle_{\mathcal{N}} \approx \pm 1$ hold approximately for $i \geq j$.

The slant of \mathbf{C} identifying the pairs (i, j) for which $c_{ij} \neq 0$, can easily be computed as follows. A classical result due to Weyl [Wey11] describes the asymptotic behavior of the Laplacian eigenvalues on manifolds. Let λ_j and λ_i denote the eigenvalues for shapes \mathcal{N} and \mathcal{M} respectively. Then, Weyl's theorem applied to 2-manifolds states that $\lambda_j \sim \frac{1}{|\mathcal{N}|} j$ and $\lambda_i \sim \frac{1}{|\mathcal{M}|} i$ as $i, j \rightarrow \infty$ (here $|\cdot|$ denotes the surface area). In other words, eigenvalues have a linear growth with the rate inversely proportional to surface area. By the previous analysis, we know that $\lambda_i \approx \lambda_j$ for $i \geq j$. Hence, it follows immediately that matrix \mathbf{C} has a slant given by $\frac{j}{i} \approx \frac{|\mathcal{N}|}{|\mathcal{M}|}$, the area ratio of the two shapes.

The following observation is crucial for the entire paper: Let $M \subset \mathcal{M}$ and $N \subset \mathcal{N}$ be two parts of the shapes \mathcal{M} and \mathcal{N} . Then, the functional correspondence between the parts M and N represented in the pair of orthogonal bases $\{\phi_i\}$ and $\{\psi_j\}$ on \mathcal{M} and \mathcal{N} , respectively, has an approximate slanted diagonal structure with the slant determined by the area ratio $|\mathcal{N}|/|\mathcal{M}|$ (see Figure 2). We emphasize that the slanted-diagonal structure of the matrix \mathbf{C} *does not* depend on the parts M and N themselves (which are typically unknown!), but only on the full shapes to which they belong.

3. Problem formulation

Let us be given a *model shape* \mathcal{M} and a collection $\{\mathcal{N}_i\}_{i=1}^p$ of p *query shapes* constituting possibly incomplete, cluttered, and non-rigidly deformed unknown parts of \mathcal{M} . Our goal is to segment \mathcal{M} into p disjoint parts $\{M_i\}$, locate the corresponding parts $\{N_i \subset \mathcal{N}_i\}$ on the input shapes, and calculate the correspondences $\tau_i : M_i \rightarrow N_i$. By *clutter* we refer to the regions $N_i^c = \mathcal{N}_i \setminus N_i$ which are redundant for achieving a full reconstruction. This may include overlaps between the \mathcal{N}_i 's, scanning artifacts, and even entire *extra parts* coming, e.g., from a different shape as we demonstrate in Figure 1. By *incompleteness* we mean that the M_i 's do not cover \mathcal{M} , i.e., there is a *missing part*

$$M_0 = \mathcal{M} \setminus \left(\bigcup_{i=1}^p M_i \right). \quad (6)$$

M_0 can be seen as clutter from the parts perspective. Figure 4 depicts our notation.

We encode the correspondences τ_i in the functional representation by the matrices \mathbf{C}_i with respect to the Laplacian eigenbasis Φ of \mathcal{M} (restricted to M_i) and the Laplacian eigenbasis Ψ_i of \mathcal{N}_i (restricted to N_i). We further assume to be given as input sets of corresponding functions on each M_i and N_i that are stacked as column vectors of (possibly differently-sized) matrices \mathbf{F}_i and \mathbf{G}_i , respectively.

Remark Since the availability of known corresponding functions is rather a restrictive assumption, in practice we avoid using this input by replacing the \mathbf{F}_i 's with a dense descriptor field \mathbf{F} calculated on \mathcal{M} (the number of columns in \mathbf{F} corresponds to the number of dimensions of the descriptor). As the \mathbf{G}_i 's, we use the descriptors




Figure 4: The notation we follow in this paper.

computed on the corresponding \mathcal{N}_i 's. A robust data fitting term accounts for descriptor mismatches.

With these premises, we formulate the simultaneous segmentation and correspondence as the following optimization problem:

$$\begin{aligned} \min_{\mathbf{C}_i, M_i \subseteq \mathcal{M}, N_i \subseteq \mathcal{N}_i} & \sum_{i=1}^p \|\mathbf{C}_i \Psi_i(N_i)^T \mathbf{G}_i - \Phi(M_i)^T \mathbf{F}_i\|_{2,1} \\ & + \lambda_{\mathcal{M}} \sum_{i=0}^p R_{\text{part}}(M_i) + \lambda_{\mathcal{N}} \sum_{i=1}^p R_{\text{part}}(N_i) \\ & + \lambda_{\text{corr}} \sum_{i=1}^p R_{\text{corr}}(\mathbf{C}_i) \\ \text{s.t. } & M_i \cap M_j = \emptyset \quad \forall i \neq j \\ & M_0 \cup M_1 \cup \dots = \mathcal{M} \\ & |M_i| = |N_i| \geq \alpha |\mathcal{N}_i|, \end{aligned} \quad (7)$$

where $\Phi(M_i)$ denotes the Laplacian eigenbasis on \mathcal{M} restricted to the part M_i , and, similarly, $\Psi_i(N_i)$ denotes the basis on \mathcal{N}_i restricted to the part N_i .

The first term in (7) is a data fitting term measuring how well the known corresponding functions are mapped between the parts and the model. The $\ell_{2,1}$ norm was chosen here to increase robustness against outliers in the input. This is especially important when one uses descriptors which are not perfectly resilient to non-rigid deformations. The second and third terms aggregating R_{part} are part regularization terms of the form $R_{\text{part}}(M) = |\partial M|$ promoting parts with short boundaries and preventing too fragmented segmentation. Note that while the regularization term applies to the missing part M_0 , the data fitting term does not. The last term aggregating R_{corr} is a regularization term imposing a prior on the correspondences themselves. Here, the prior comes in the form of a penalty promoting the slanted diagonal structure of each \mathbf{C}_i with the slant proportional to the ratio $|\mathcal{N}_i|/|\mathcal{M}|$ as detailed in the sequel.

Finally, the set of constraints renders the problem a proper seg-

mentation task, enforcing a complete covering and exclusivity of the segments M_i . The area constraint enforces the non-cluttered matching areas N_i to be equal. For cases where there exist both clutter in the parts and missing elements, we introduce the inequality term putting a lower bound on the part areas to avoid the trivial solution. In such cases, one has to impose a prior on the resulting non-cluttered area being greater than some percentage α of the entire cluttered part.

Since problem (7) is intractable in its combinatorial formulation, we consider a relaxation of the parts to continuous membership functions $u_i : \mathcal{M} \rightarrow [0, 1]$ to encode the M_i 's, and $v_i : \mathcal{N}_i \rightarrow [0, 1]$ to encode the N_i 's. Assuming that \mathcal{M} is discretized with m vertices, and each \mathcal{N}_i is discretized with n_i vertices, the relaxed and discretized optimization problem can be summarized as

$$\begin{aligned} \min_{\mathbf{C}_i, \mathbf{u}_i, \mathbf{v}_i} & \sum_{i=1}^p \|\mathbf{C}_i \mathbf{A}_i(\boldsymbol{\eta}(\mathbf{u}_i)) - \mathbf{B}_i(\boldsymbol{\eta}(\mathbf{v}_i))\|_{2,1} + \lambda_{\mathcal{M}} \sum_{i=0}^p R_{\text{part}}(\boldsymbol{\eta}(\mathbf{u}_i)) \\ & + \lambda_{\mathcal{N}} \sum_{i=1}^p R_{\text{part}}(\boldsymbol{\eta}(\mathbf{v}_i)) + \lambda_{\text{corr}} \sum_{i=1}^p R_{\text{corr}}(\mathbf{C}_i) \\ \text{s.t.} & \sum_{i=0}^p \boldsymbol{\eta}(\mathbf{u}_i) = \mathbf{1} \\ & \mathbf{a}_{\mathcal{M}}^T \boldsymbol{\eta}(\mathbf{u}_i) = \mathbf{a}_{\mathcal{N}_i}^T \boldsymbol{\eta}(\mathbf{v}_i) \geq \alpha \mathbf{a}_{\mathcal{N}_i}^T \mathbf{1} \end{aligned} \quad (8)$$

Note that, differently from [RCB*16], here we solve p matching problems simultaneously (one per part) under covering and exclusivity constraints.

In the problem above, \mathbf{a} denote the vectors of discrete area elements on the corresponding shapes, and $\boldsymbol{\eta}(t) = \frac{1}{2} \tanh\left(6(t - \frac{1}{2})\right) + \frac{1}{2}$ is an element-wise non-linear transformation used to restrict the indicators at each vertex to the range $[0, 1]$ (see inset). Function $\boldsymbol{\eta}(t)$ was chosen according to [RCB*16]. The matrices $\mathbf{A}_i(\boldsymbol{\eta}(\mathbf{u}_i)) = \boldsymbol{\Phi}^T \text{diag}(\mathbf{u}_i) \mathbf{F}_i$ and $\mathbf{B}_i(\boldsymbol{\eta}(\mathbf{v}_i)) = \boldsymbol{\Psi}_i^T \text{diag}(\mathbf{v}_i) \mathbf{G}_i$ denote the representation coefficients of the input descriptor fields restricted to their respective parts.

As the regularization term of the segments we use a discretized version of the intrinsic Mumford-Shah functional introduced in [BB08a]

$$R_{\text{part}}(\boldsymbol{\eta}(\mathbf{u})) = \int_{\mathcal{M}} \xi(\boldsymbol{\eta}(\mathbf{u})) \|\nabla_{\mathcal{M}} \boldsymbol{\eta}(\mathbf{u})\| da \approx \mathbf{a}_{\mathcal{M}}^T \mathbf{g}, \quad (9)$$

where $\xi(t) \approx \delta\left(t - \frac{1}{2}\right)$, and the vector \mathbf{g} contains as its elements the values of the discretized intrinsic gradient norm of $\boldsymbol{\eta}(\mathbf{u})$ computed on the tangent bundle of \mathcal{M} , weighted element-wise by $\xi(\boldsymbol{\eta}(\mathbf{u}))$.

For the regularization of the functional maps \mathbf{C}_i , we follow [RCB*16],

$$\begin{aligned} R_{\text{corr}}(\mathbf{C}) &= \|\mathbf{C} \odot \mathbf{W}\|_{\mathbf{F}}^2 + \lambda_1 \sum_{i \neq j} (\mathbf{C}^T \mathbf{C})_{ij}^2 \\ &+ \lambda_2 \sum_i ((\mathbf{C}^T \mathbf{C})_{ii} - d_i)^2. \end{aligned} \quad (10)$$




Figure 5: An example showing the decrease in cost during the alternating minimization of the different sets of variables, \mathbf{C}_i (red), \mathbf{v}_i (green) and \mathbf{u}_i (blue).

Here, the first term containing an element-wise product of \mathbf{C} with the funnel-shaped weight matrix \mathbf{W} promotes the slanted-diagonal structure of \mathbf{C} . The elements of the weight matrix are given by

$$w_{mn} = e^{-\sigma \sqrt{m^2 + n^2}} \left\| \frac{\mathbf{n}}{\|\mathbf{n}\|} \times ((m, n)^T - \mathbf{p}) \right\|. \quad (11)$$

The slanted diagonal of \mathbf{W} is a line segment $\ell(t) = \mathbf{p} + t \frac{\mathbf{n}}{\|\mathbf{n}\|}$ with $t \in \mathbb{R}$, where $\mathbf{p} = (1, 1)^T$ is the matrix origin, and $\mathbf{n} = (1, |\mathcal{N}_i|/|\mathcal{M}|)^T$ is the line direction with slope $|\mathcal{N}_i|/|\mathcal{M}|$. The second factor in w_{mn} is the distance from the slanted diagonal ℓ , and $\sigma > 0$ regulates the spread around ℓ . In our experiments we set $\sigma = 0.03$. The second term in $R_{\text{corr}}(\mathbf{C})$ promotes orthogonality of \mathbf{C} (area-preserving maps), while the third term regularizes its rank by setting $\{d_i\}_{i=1}^r = 1$ and $\{d_i\}_{i=r+1}^k = 0$. Following [RCB*16], the estimation of r is done by

$$r = \max \{i \mid \lambda_i^{\mathcal{N}} < \max_j \lambda_j^{\mathcal{M}}\}. \quad (12)$$

4. Implementation

We solve the optimization problem by means of a threefold alternating minimization. To this end, we make use of the masks \mathbf{W}_i to initialize the matrices \mathbf{C}_i by applying the transformation $\mathbf{C}_i = \mathbf{1}\mathbf{1}^T - \frac{\mathbf{W}_i}{\max(\mathbf{W}_i)}$. We then minimize over the partial functional maps \mathbf{C}_i , the model indicator functions \mathbf{u}_i and the parts indicator functions \mathbf{v}_i in a cyclic manner, keeping the other parameters fixed (see Algorithm 1).

Although this algorithm is not guaranteed to converge, in practice we observed a strictly decreasing cost value as the one shown in Figure 5. For the different minimization steps we used the conjugate gradient solver supplied by the Manopt toolbox [BMAS14]. Since this solver does not support constraints inherently, they were replaced by large quadratic penalties. In order to further refine the solution for the functional maps \mathbf{C}_i we added a k -dimensional ICP step [RCB*16]. We noticed this step helps especially when the descriptors are performing poorly. The parameters were changed according to the required setting of the experiment. For instance, in




Figure 6: Comparison between partial functional maps [RCB*16] and our method in a “perfect” puzzle setting. For each method we show the membership functions of each part with respect to the model (first two rows), and the color-coded correspondence between parts and model (last two rows). For PFM, the fact that each part is matched independently leads to different parts covering overlapping areas on the model (see, e.g., the four legs). This ambiguity is completely resolved by our method as all parts are matched jointly to the template, yielding a regularizing effect on the correspondence.

the non-isometric experiment (Figure 7) we set $\lambda_1 = 0$ to allow changes of areas.

Input : model \mathcal{M} , parts $\{\mathcal{N}_i\}_{i=1}^p$

Output: segments $\{\mathbf{u}_i\}$, $\{\mathbf{v}_i\}$ and maps $\{\mathbf{C}_i\}$

initialization: $\mathbf{C}_i = \mathbf{1}\mathbf{1}^\top - \frac{\mathbf{W}_i}{\max(\mathbf{W}_i)}$, $\mathbf{u}_i = \mathbf{1}$, $\mathbf{v}_i = \mathbf{1}$

while decrease in energy $> \epsilon$ **do**

 fix $\{\mathbf{u}_i\}$ and $\{\mathbf{v}_i\}$ in Eq. (8) and solve for $\{\mathbf{C}_i\}$;

 run spectral ICP [RCB*16] for all \mathbf{C}_i ;

 fix $\{\mathbf{C}_i\}$ and $\{\mathbf{u}_i\}$ in Eq. (8) and solve for $\{\mathbf{v}_i\}$;

 fix $\{\mathbf{C}_i\}$ and $\{\mathbf{v}_i\}$ in Eq. (8) and solve for $\{\mathbf{u}_i\}$;

end

Algorithm 1: Our pipeline for solving non-rigid puzzles.

5. Experimental results

Our method was implemented in C++/Matlab, and executed on an Intel i7-4710MQ 2.50GHz CPU with 8 logical cores. Typical running times for matching 5 parts to a template of about 10K vertices were 20 minutes (end-to-end).

Data. In our experiments we use both synthetic and real data. The synthetic dataset is made up of shapes from the TOSCA [BBK08] and FAUST [BRLB14] benchmarks. In order to avoid compatible meshings and make the dataset more realistic, each TOSCA model is independently remeshed to $\sim 10K$ vertices by iterative pair contractions [GH97]. All FAUST templates are kept at their original resolution ($\sim 7K$). The second dataset is composed of real scans acquired with a calibrated Asus Xtion Pro Live RGB-D sensor and then fused into a dense 3D model (about 30K vertices) by DVO-SLAM [KSC13].

The shapes from these datasets are decomposed into a control-

lable amount of parts by Voronoi decomposition and consensus segmentation [RBC14]; the former approach leads to generic surface patches having similar area, while the latter tends to produce more semantically meaningful parts (e.g., arms and feet).

Features. Unless differently stated, as dense descriptor fields for the data term in (7) we use 350-dimensional SHOT signatures [TSDS10]. These are rotation-invariant local features with *no* isometry invariance, but whose locality properties result in a higher resilience towards boundary effects than classical spectral features [SOG09, ASC11]. Note that we compute dense descriptors for all shape points, including those lying along the boundaries.




Figure 7: Comparison between our method (top row) and PFM (bottom row) on real data. The parts shown on the right were acquired with a 3D scanner.




Figure 8: An example of our matching pipeline when dealing with overlapping parts. The optimization process alternates between the membership functions on the model (top row) and those on the parts (bottom row). At each alternating step, the membership functions are optimized jointly on the respective parts. Note that in this example there is more than one possible solution (e.g., the blue parts are redundant).

5.1. Perfect puzzle

In Figure 6 we show an example of a solution obtained with our method in a basic setting. The input data are five non-overlapping pieces taken from nearly isometric deformations of the model, forming a covering set of the model. SHOT descriptors were used in the data fitting term. No additional clutter is introduced. For this experiment, we compare with the partial functional maps (PFM) method of Rodolà *et al.* [RCB*16] applied to each part separately, resulting in five independent PFM matching problems (one per part).

We performed a similar comparison with real data acquired by a 3D sensor. For this experiment we use the upper part of a shape from FAUST as a template, and portions of a real scanning as the data. Differently from the previous experiment where dense SHOT descriptors are used, here we employ Gaussians supported at ~ 15 hand-picked matches as data features. The results are reported in Figure 7.

5.2. Overlapping pieces

A more interesting setup is obtained when allowing the different pieces to have non-zero overlap, as illustrated in Figure 8. In Figure 9 we show additional results obtained in this setting. To make the experiment even more challenging, we produce the input parts by decomposing into five components two different *non-isometric* shapes from the FAUST dataset. The decomposition is performed so as to allow large areas of overlap between the pieces. We see that our method copes well with both sources of nuisance even if these show up simultaneously: Overlapping areas are correctly segmented, while the lack of isometry does not have a significant impact on the quality of the correspondence.

5.3. Incomplete and noisy data

In practical situations, it may happen that the parts at our disposal do not provide a complete covering of the template model. As described in Section 3, our method naturally allows handling scenarios where some of the parts are missing. This is simply done by introducing a lower bound on the part areas, reflecting some prior knowledge on the amount of missing area; note that, in the absence of clutter, this is directly given by the difference of template area and the sum of the parts. In practice we implement this by defining a membership function to represent the missing part, which is then treated the same way as the others (*i.e.*, we demand regularity on the missing area, yet provide no data term).

In Figure 1 we show an example of such a scenario, with additional ‘extra’ pieces that do not belong to the model (the head of the cat). In this noisy setting, the outlier shape is automatically excluded from the final solution due to a lack of mutual support with the rest of the data. Another example of this challenging scenario is given in Figure 10.

6. Discussion and conclusions

In this paper we introduced a method for solving 3D non-rigid puzzle problems. We formulated the problem as one of partial functional correspondence among an input set of surface pieces and a full template model known in advance. The pieces are matched to the template in a joint fashion, and an optimization process alternates between optimizing for the dense part-to-whole correspondence and the segmentation of the model. We showed how the set of constraints imposed on the plurality of the pieces has a regularizing effect on the solution, leading to accurate part alignment even in challenging scenarios. The framework we presented is remarkably flexible, and can be easily adapted to deal with missing or overlapping pieces, moderate amounts of clutter, and outliers.




Figure 9: Non-rigid puzzle alignment between overlapping parts. Shown are the final segmentation obtained by our method (left), the dense matchings between the parts and the model (middle), and the normalized geodesic error (shown as a heatmap) to the ground-truth correspondence (right). With the exception of the final column, corresponding points have the same color whereas white color denotes no match. Despite the lack of isometry (two different individuals) and the large overlap, our method correctly identifies non-overlapping subregions on all the parts, providing a perfect covering of the template. Note that this is not the only possible solution, as the optimization problem we solve may have multiple optima.




Figure 10: In this example, an additional outlier piece (the human arm) is included in the input set. Our method treats extra pieces as clutter; the arm is automatically selected by the matching process, and completely excluded from the final solution. Note how the presence of the extraneous part did not affect the quality of the correspondence.

Limitations. One of the main limitations, which our method inherits from the functional maps framework, is the need for a reasonably good data term, implying that one has to provide some corresponding functions between the model and the query shapes. This is especially important when the shapes being matched are not nearly isometric. In order to allow a fully-automatic pipeline, dense descriptors are used as such corresponding functions. Yet, in real-world settings when the data is contaminated by noise and scanning artifacts, obtaining invariant descriptors is a major challenge.

Acknowledgments

We thank Michael Moeller for fruitful discussions and Christian Kerl for the technical support in the real scanning experiments. AB and OL are supported by the ERC Starting Grant No. 335491. ER and MB are supported by the ERC Starting Grant No. 307047. DC is supported by the ERC Consolidator Grant “3D Reloaded”.

References

- [AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust pairwise surface registration. *TOG* 27, 3 (2008), 85. 2
- [ART15] ALBARELLI A., RODOLÀ E., TORSELLO A.: Fast and accurate surface alignment through an isometry-enforcing game. *Pattern Recognition* 48 (2015), 2209–2226. 2
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops* (2011). 6
- [BB08a] BRONSTEIN A. M., BRONSTEIN M. M.: Not only size matters: regularized partial matching of nonrigid shapes. In *Proc. NORDIA* (2008). 2, 5
- [BB08b] BRONSTEIN A. M., BRONSTEIN M. M.: Regularized partial matching of rigid shapes. In *Proc. ECCV*. 2008. 2
- [BBB09] BRONSTEIN A., BRONSTEIN M., BRUCKSTEIN A., KIMMEL R.: Partial similarity of objects, or how to compare a centaur to a horse. *IJCV* 84, 2 (2009), 163–183. 2
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 2
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 6
- [BCBB15] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer Graphics Forum* (2015). 2
- [BMAS14] BOUMAL N., MISHRA B., ABSIL P.-A., SEPULCHRE R.: Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research* 15 (2014), 1455–1459. URL: <http://www.manopt.org>. 5
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (June 2014). 6
- [BW*14] BRUNTON A., WAND M., WUHRER S., SEIDEL H.-P., WEINKAUF T.: A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models* 76, 2 (2014), 70–85. 2
- [CRA*16] COSMO L., RODOLÀ E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* (2016). 2
- [CRB*16] COSMO L., RODOLÀ E., BRONSTEIN M. M., TORSELLO A., CREMERS D., SAHILLIOĞLU Y.: Shrec’16: Partial matching of deformable shapes. In *Proc. 3DOR* (2016). 2
- [DTF*15] DOU M., TAYLOR J., FUCHS H., FITZGIBBON A., IZADI S.:

- 3D scanning deformable objects with a single RGBD sensor. In *Proc. CVPR* (2015). 2
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216. 6
- [HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. *TOG* 25, 3 (2006), 569–578. 2
- [HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 177–186. 2
- [HWG14] HUANG Q., WANG F., GUIBAS L. J.: Functional map networks for analyzing and exploring large shape collections. *TOG* 33, 4 (2014), 36. 2
- [KBB*13] KOVNATSKY A., BRONSTEIN M., BRONSTEIN A., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Comput. Graph. Forum* 32, 2pt4 (2013), 439–448. 3
- [KBBV14] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. *CoRR abs/1412.8070* (2014). 2
- [KSC13] KERL C., STURM J., CREMERS D.: Dense visual slam for rgb-d cameras. In *Proc. IROS* (2013). 6
- [LBB12] LITANY O., BRONSTEIN A. M., BRONSTEIN M. M.: Putting the pieces together: Regularized multi-part shape matching. In *Proc. NORDIA* (2012). 2
- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. In *Proc. SGP* (2008), pp. 1421–1430. 2
- [MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDERGHEYNST P.: Geodesic convolutional neural networks on riemannian manifolds. In *Proc. 3dRR* (2015). 2
- [MS89] MUMFORD D., SHAH J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure and Applied Math.* 42, 5 (1989), 577–685. 2
- [NFS15] NEWCOMBE R. A., FOX D., SEITZ S. M.: DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proc. CVPR* (2015). 2
- [NIH*11] NEWCOMBE R. A., IZADI S., HILLIGES O., MOLYNEAUX D., KIM D., DAVISON A. J., KOHI P., SHOTTON J., HODGES S., FITZGIBBON A.: Kinectfusion: Real-time dense surface mapping and tracking. In *Proc. ISMAR* (2011), pp. 127–136. 2
- [OBBS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July 2012), 30:1–30:11. 2, 3
- [PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum* 32, 2pt4 (2013), 459–468. 3
- [RBA*12] RODOLÀ E., BRONSTEIN A., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *Proc. CVPR* (June 2012), pp. 182–189. 2
- [RBC14] RODOLÀ E., BULÒ S. R., CREMERS D.: Robust region detection via consensus segmentation of deformable shapes. *Computer Graphics Forum* 33, 5 (2014), 97–106. 6
- [RCB*16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* (2016). 2, 3, 5, 6, 7
- [RTH*13] RODOLÀ E., TORSELLO A., HARADA T., KUNIYOSHI Y., CREMERS D.: Elastic net constraints for shape matching. In *Proc. ICCV* (December 2013), pp. 1169–1176. 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP* (2009). 6
- [SY14a] SAHILLIOĞLU Y., YEMEZ Y.: Multiple shape correspondence by dynamic programming. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 121–130. 2
- [SY14b] SAHILLIOĞLU Y., YEMEZ Y.: Partial 3-d correspondence from shape extremities. *Computer Graphics Forum* 33, 6 (2014), 63–76. 2
- [TRA11] TORSELLO A., RODOLÀ E., ALBARELLI A.: Multiview registration via graph diffusion of dual quaternions. In *Proc. CVPR* (2011). 2
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proc. ECCV* (2010), pp. 356–369. 6
- [VC02] VESE L. A., CHAN T. F.: A multiphase level set framework for image segmentation using the Mumford and Shah model. *IJCV* 50, 3 (2002), 271–293. 2
- [vKZH13] VAN KAICK O., ZHANG H., HAMARNEH G.: Bilateral maps for partial matching. *Computer Graphics Forum* 32, 6 (2013), 189–200. 2
- [vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 1, 2
- [Wey11] WEYL H.: Über die asymptotische Verteilung der Eigenwerte. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* (1911), 110–117. 4
- [WHC*15] WEI L., HUANG Q., CEYLAN D., VOUGA E., LI H.: Dense human body correspondences using convolutional networks. *arXiv 1511.05904* (2015). 2
- [WSSC11] WINDHEUSER T., SCHLICKWEI U., SCHMIDT F. R., CREMERS D.: Large-scale integer linear programming for orientation preserving 3d shape matching. *Computer Graphics Forum* 30, 5 (2011), 1471–1480. 2