Contents

1 Partial Single- and Multi-Shape Dense Correspondence Using Functional Maps **1**
1.1 Introduction . . . 1
1.1.1 Related work . . . 1
1.2 Full non-rigid shape correspondence . . . 3
1.2.1 Functional representation . . . 3
1.2.2 Joint diagonalization . . . 4
1.3 Partial functional correspondence . . . 5
1.4 Multi-part partial functional maps . . . 7
1.4.1 Non-rigid puzzles . . . 8
1.4.2 Implementation . . . 11
1.4.3 Experimental results . . . 11
1.4.4 Discussion and conclusions . . . 13
1.5 Fully-spectral partial functional maps . . . 14
1.5.1 Implementation . . . 17
1.5.2 Experimental results . . . 17
1.5.3 Discussion and conclusions . . . 19

Chapter 1

Partial Single- and Multi-Shape Dense Correspondence Using Functional Maps

1.1 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 [56]. 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 than 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). In this chapter we will deal with describe recent methods that handle all these challenging settings.

1.1.1 Related work

Albeit one of the most broadly studied problems in the domain of geometry processing, correspondence is far from being solved, especially in some challenging settings. We refer the reader to recent survey papers [4, 56] 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 [1, 2]. Bronstein *et al.* [10] used a regularized ICP approach where the matching parts are explicitly modeled, and proposed a functional similar to the Mumford-Shah [38, 57] imposing part regularity. Litany *et al.* [32] 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 [11]. Bronstein *et al.* [7,9] combined metric distortion minimization with optimization over regular matching parts. Rodolà *et al.* [44, 49] relaxed the regularity requirement by allowing sparse correspondences. Windheuser *et al.* [59] proposed an integer linear programming solution for dense elastic matching. Sahillioğlu and Yemez [50] 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 [12] and the design of robust descriptors for partial matching [55]; in the context of shape collections, partial correspondence has been considered in [15, 22]. Several works tried to employ machine learning methods to deal with partial matches. Masci *et al.* [35] 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.* [58] 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 [31] had limited success due to sensitivity to initialization and to the underlying assumption of small deformations. Recent works of [19, 39] generalizing the Kinect fusion approach [40], 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.* [41] introduced *functional maps*, representing correspondences between functional spaces on the respective shapes. While not intended for partial correspondence, followup works [28] showed that functional maps and similar constructions can handle certain settings with missing parts. Rodolà *et al.* [46] introduced *partial functional correspondence*, an extension of [41] where matched parts are explicitly modeled and regularized in a manner similar to [7, 9]. This method has achieved state-of-the-art performance on the recent SHREC'16 Partial Correspondence benchmark [16].

Multiple shape correspondence in the rigid settings has been addressed in numerous works, including [21, 32]. In the non-rigid setting, functional maps for large shape collections have been explored in [22].

The rest of the chapter is organized as follows. Section 1.2 formulates the problem of dense correspondence between two full shapes, and describes the basic functional maps framework to solve it. The *Joint diagonalization* method is further presented, which allows handling a deviation from the isomtry assumption, and sets the ground for upcoming sections. Section 1.3 extends the matching problem to a partial setting, namely, when one shape mathces only a subset of the other. An extension to multiple parts and possible addition of clutter is described in Section 1.4. The partial matching algorithm used in Sections 1.3, 1.4 suffers from poor scalability to large number of points due to its spatial component. In Section 1.5 a fully-spectral approach is described, that achieves similar results at a much lower computational complexity.

1.2 Full non-rigid shape correspondence

We model shapes as 2-manifolds $\mathcal{M}$ (possibly with boundary $\partial\mathcal{M}$) equipped with the area element $\mathrm{d}\mu$ induced by the standard metric. The intrinsic gradient $\nabla_{\mathcal{M}}$ and the positive semi-definite Laplace-Beltrami operator $\Delta_{\mathcal{M}}$ generalize the corresponding notions from flat spaces to manifolds. The Laplacian admits an eigen-decomposition

$$\Delta_{\mathcal{M}}\phi_i(x) = \lambda_i\phi_i(x) \qquad x \in \mathrm{int}(\mathcal{M}) \tag{1.1}$$

$$\langle \nabla_{\mathcal{M}}\phi_i(x), \hat{n}(x)\rangle = 0 \qquad x \in \partial\mathcal{M}, \tag{1.2}$$

with Neumann boundary conditions (1.2), where $\hat{n}$ is the normal vector to the boundary. Here, $0 = \lambda_1 \leq \lambda_2 \leq \ldots$ are eigenvalues and $\phi_1, \phi_2, \ldots$ are the corresponding eigenfunctions. Due to the isometry invariance of the Laplacian, nearly-isometric shapes will have approximately the same eigenvalues and eigenspaces (up to orthogonal transformation).

By analogy to the Euclidean case, the Laplace operator $\Delta_{\mathcal{M}}$ allows us to extend Fourier analysis to manifolds. Since the eigenfunctions of the Laplacian form an orthonormal basis of $L^2(\mathcal{M}) = \{f : \mathcal{M} \to \mathbb{R} \mid \int_{\mathcal{M}} f^2 \mathrm{d}\mu < \infty\}$, the space of square-integrable functions on $\mathcal{M}$, any function $f \in L^2(\mathcal{M})$ can be represented via the Fourier series expansion

$$f(x) \;=\; \sum_{i\geq 1} \langle f, \phi_i\rangle_{\mathcal{M}}\phi_i(x)\,, \tag{1.3}$$

where we use the standard $L^2(\mathcal{M})$ inner product defined as $\langle f, g\rangle_{\mathcal{M}} = \int_{\mathcal{M}} fg \; \mathrm{d}\mu$.

1.2.1 Functional representation

The functional maps framework was proposed by Ovsjanikov *et al.* [41] in 2012. The main idea is to identify correspondences between shapes by a linear operator $T : L^2(\mathcal{M}) \to L^2(\mathcal{N})$, mapping functions on $\mathcal{M}$ to functions on $\mathcal{N}$. One can easily see that classical point-to-point correspondences constitute a special case where delta functions are mapped to delta functions.

As a linear operator, T admits a matrix representation $\mathbf{C} = (c_{ij})$ with coefficients computed as follows. Let $\{\phi_i\}_{i\geq 1}$ and $\{\psi_j\}_{j\geq 1}$ be orthonormal bases on $L^2(\mathcal{M})$ and $L^2(\mathcal{N})$, respectively, and let $f \in L^2(\mathcal{M})$. Then, the action of T on f can be written as

$$\begin{aligned} Tf \;&=\; T\sum_{i\geq 1}\langle f, \phi_i\rangle_{\mathcal{M}}\phi_i = \sum_{i\geq 1}\langle f, \phi_i\rangle_{\mathcal{M}} T\phi_i \\ &=\; \sum_{ij\geq 1}\langle f, \phi_i\rangle_{\mathcal{M}} \underbrace{\langle T\phi_i, \psi_j\rangle_{\mathcal{N}}}_{c_{ji}} \psi_j\,. \end{aligned} \tag{1.4}$$

By choosing as functional bases $\{\phi_i\}_{i\geq 1}$, $\{\psi_j\}_{j\geq 1}$ the Laplacian eigenfunctions on the respective manifolds, one obtains a particularly compact representation for the functional map: this choice allows to truncate the series (1.4) after the first k terms as a band-limited approximation of the original map, by analogy with Fourier analysis. This results in a $k \times k$ matrix $\mathbf{C}$ encoding the functional correspondence, where k is typically chosen to be a small number (20 to 100 in practice). If, in addition, the functional map T is built on top of a near-isometry, one obtains $c_{ji} = \langle T\phi_i, \psi_j\rangle_{\mathcal{N}} \approx \pm\delta_{ji}$ since near-isometric shapes have corresponding eigenfunctions (up to sign in case of simple spectra). The resulting matrix $\mathbf{C}$ thus manifests a diagonally dominant structure.

Now assume to be given q corresponding functions $g_i \approx Tf_i$, $i = 1, \ldots, q$ and let $\mathbf{A} = (\langle \phi_i, f_j \rangle_{\mathcal{M}})$ and $\mathbf{B} = (\langle \psi_i, g_j \rangle_{\mathcal{N}})$ be the $k \times q$ matrices of Fourier coefficients of the given corresponding functions. The functional correspondence problem considered in [41] has the general form of

$$\min_{\mathbf{C}} \|\mathbf{CA} - \mathbf{B}\|_{\mathrm{F}}^2 , \tag{1.5}$$

with the additional orthogonality constraint $\mathbf{C}^\top \mathbf{C} = \mathbf{I}$ if the underlying map is known to be area-preserving [41].

A recent paradigm shift in the shape matching problem was introduced by Ovsjanikov *et al.* [41]. The authors proposed to model correspondences among two shapes by means of a linear operator $T : L^2(\mathcal{M}) \to L^2(\mathcal{N})$, 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 $\{\psi_i\}_{i \geq 1}$ on $L^2(\mathcal{M})$ and $L^2(\mathcal{N})$, respectively, and let us fix some function $f \in L^2(\mathcal{M})$. Then

$$\begin{aligned} Tf &= T \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} T\phi_i \\ &= \sum_{ij \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \underbrace{\langle T\phi_i, \psi_j \rangle_{\mathcal{N}}}_{c_{ij}} \psi_j . \end{aligned} \tag{1.6}$$

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

Choosing as the bases the eigenfunctions $\{\phi_i\}_{i \geq 1}$, $\{\psi_i\}_{i \geq 1}$ of the respective Laplacians on the two shapes yields a particularly convenient representation for the functional map [41]. By analogy with Fourier analysis, this choice allows to truncate the series (1.6) 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 [26, 43] as a prior for shape matching problems.

1.2.2 Joint diagonalization

When dealing with *non*-isometric shapes, the diagonally dominant structure of $\mathbf{C}$ is broken since the approximate equality $c_{ji} = \langle T\phi_i, \psi_j \rangle_{\mathcal{N}} \approx \pm\delta_{ji}$ ceases to hold. In [29] it was proposed to find a pair of new bases $\{\hat{\phi}_i, \hat{\psi}_i\}_{i=1}^k$ in which $\mathbf{C}$ still has a near-diagonal structure. The new bases are constructed as linear combinations of k standard Laplacian eigenfunctions,

$$\hat{\phi}_i = \sum_{j=1}^{k} p_{ji} \phi_j , \qquad \hat{\psi}_i = \sum_{j=1}^{k} q_{ji} \psi_j \tag{1.7}$$

where $\mathbf{P}, \mathbf{Q}$ are the $k \times k$ matrices with the combination coefficients. It is easy to check that the requirement for orthogonality of the new bases $\langle \hat{\phi}_i, \hat{\phi}_j \rangle_{\mathcal{M}} = \delta_{ij}$ and $\langle \hat{\psi}_i, \hat{\psi}_j \rangle_{\mathcal{N}} = \delta_{ij}$ implies the orthogonality of the matrices $\mathbf{P}^\top \mathbf{P} = \mathbf{I}$ and $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$. Further, the coefficients

of $\{f_i, g_i\}$ in the new bases can be expressed as $\hat{\mathbf{A}} = \mathbf{P}^\top \mathbf{A}$ and $\hat{\mathbf{B}} = \mathbf{Q}^\top \mathbf{B}$. The goal is to find matrices $\mathbf{P}, \mathbf{Q}$ resulting in "quasi-harmonic" bases $\{\hat{\phi}_i, \hat{\psi}_i\}$, *i.e.*, that behave approximately as eigenfunctions of the Laplacian, while being *coupled* in the sense $\hat{\mathbf{A}} \approx \hat{\mathbf{B}}$. Due to the coupling, the new basis functions behave consistently resulting in almost perfectly diagonal $\mathbf{C}$ even in the absence of a perfect isometry.

The orthogonal basis $\{\hat{\phi}_i\}$ behaves as the eigenbasis of $\Delta_\mathcal{M}$ if it minimizes the Dirichlet energy $\sum_{i=1}^k \langle \hat{\phi}_i, \Delta_\mathcal{M} \hat{\phi}_i \rangle_\mathcal{M} = \operatorname{tr}(\mathbf{P}^\top \mathbf{\Lambda}_\mathcal{M} \mathbf{P})$, where $\mathbf{\Lambda}_\mathcal{M}$is a diagonal matrix of the first k eigenvalues of $\Delta_\mathcal{M}$, and where we used the fact that $\langle \phi_i, \Delta_\mathcal{M} \phi_j \rangle_\mathcal{M} = \lambda_j \delta_{ij}$. Alternatively, the trace term can be replaced by an off-diagonal penalty [13], arriving at the optimization problem

$$\begin{aligned} \min_{\mathbf{P},\mathbf{Q}} \quad & \operatorname{off}(\mathbf{P}^\top \mathbf{\Lambda}_\mathcal{M} \mathbf{P}) + \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_\mathcal{N} \mathbf{Q}) + \mu \|\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_\mathrm{F}^2 \\ \text{s.t.} \quad & \mathbf{P}^\top \mathbf{P} = \mathbf{I}, \quad \mathbf{Q}^\top \mathbf{Q} = \mathbf{I}, \end{aligned} \tag{1.8}$$

where $\operatorname{off}(\mathbf{A}) = \sum_{i \neq j} a_{ij}^2$. Problem (1.8) can be interpreted as a *joint approximate diagonalization* of the Laplacians $\Delta_\mathcal{M}$ and $\Delta_\mathcal{N}$ [29]. Note that if $\mu = 0$ (*i.e.*, no coupling) the global solution to (1.8) is $\mathbf{P} = \mathbf{Q} = \mathbf{I}$, resulting in the standard eigenfunctions of $\Delta_\mathcal{M}$ and $\Delta_\mathcal{N}$ when plugged into (1.7).

The orthogonal matrices $\mathbf{P}$ and $\mathbf{Q}$ act as rotations and reflections of the original eigenbases, trying to align them in the k-dimensional eigenspace. Because of this interpretation, it is possible to simplify problem (1.8) by optimizing for a new basis on one shape only and keeping the other fixed to the standard Laplacian eigenfunctions,

$$\min_{\mathbf{Q} \in S(k,k)} \quad \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_\mathcal{N} \mathbf{Q}) + \mu \|\mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_\mathrm{F}^2, \tag{1.9}$$

where $S(n,k) = \{\mathbf{X} \in \mathbb{R}^{n \times k} : \mathbf{X}^\top \mathbf{X} = \mathbf{I}_k\}$ denotes the *Stiefel manifold* of $n \times k$ orthogonal matrices (when $k < n$, such matrices are also called *ortho-projections*). Problems (1.8–1.9) are instances of *manifold optimization* and can be solved using efficient numerical techniques performing optimization on the matrix manifold [6].

Robust formulation. In practical settings, the corresponding functions f_i, g_i might be noisy, such that $T f_i \neq g_i$ for some i's. As a result, some of the columns in the data term $\mathbf{A} - \mathbf{Q}^\top \mathbf{B}$ might have large norm. A standard way to cope with such outliers is to replace the ℓ_2 (Frobenius) norm in (1.9) with a robust matrix norm $\|\mathbf{X}\|_{2,1} = \sum_i \|\mathbf{x}^i\|_2$ promoting column-wise sparsity (here $\mathbf{x}^i$ is the ith column of $\mathbf{X}$). When the input functions are different dimensions of a high-dimensional descriptor field, this has the effect of discarding entire feature channels from the data. Note that robustness to point (as opposed to channel) mismatches may be achieved by row-wise sparsity in the *spatial* domain, however doing so would sacrifice the benefits of shifting to a spectral representation. Finally, the presence of the $\ell_{2,1}$ norm makes the objective function non-smooth. In such a setting, non-smooth manifold optimization techniques such as MADMM [27] can be employed to reach a good local optimum.

1.3 Partial functional correspondence

Assume now to be given a full shape $\mathcal{N}$ and a *partial* shape $\mathcal{M}$ that is approximately isometric to some (unknown) sub-region $\mathcal{N}' \subseteq \mathcal{N}$. We are interested in determining a *partial functional map* $T : L^2(\mathcal{M}) \to L^2(\mathcal{N})$ mapping functions on $\mathcal{M}$ to functions supported on the region $\mathcal{N}'$.

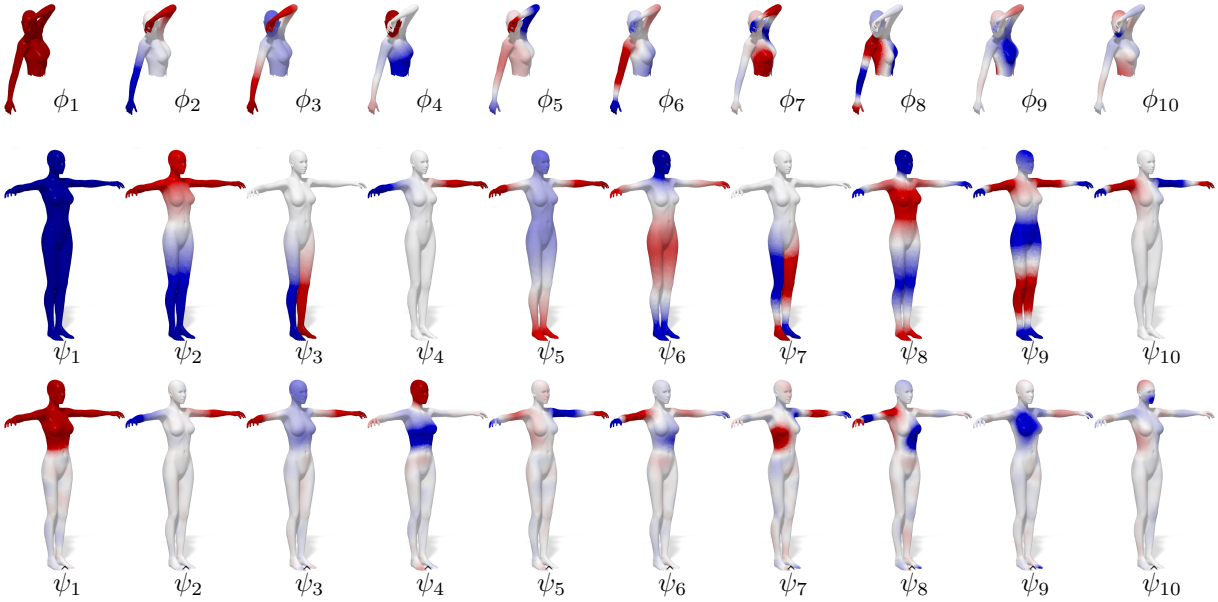

Figure 1.1: The standard Laplacian eigenfunctions (two first rows) are strongly affected by the lack of perfect isometry and in the presence of missing parts. In the top and middle rows we show the first ten eigenfunctions $\{\phi_i\}_{i=1}^{10}$ and $\{\psi_j\}_{j=1}^{10}$ on a partial and full shape respectively; note the inconsistent behavior at corresponding indices. In the bottom row we show the optimal basis functions $\{\hat{\psi}_j\}_{j=1}^{10}$ obtained with the method presented in Section 1.5: the new basis manifests the same behavior as in the first row, and is at the same time localized on the latent corresponding part.

Recently, Rodolà *et al.* [46] showed that for each "partial" eigenfunction ϕ_j (*i.e.*, each eigenfunction of the part $\mathcal{M}$) there exists a corresponding "full" eigenfunction ψ_i of $\mathcal{N}$ for some $i \geq j$ (see for example ϕ_3 and ψ_5 in Figure 1.1). Differently from the full-to-full setting, where the correspondence is observed for $i = j$, here the inequality $i \geq j$ induces a *slanted-diagonal* structure on matrix $\mathbf{C}$. In particular, under the correct isometry encoded in T (*i.e.*, the image $T\phi_j$ is localized to $\mathcal{N}' \subseteq \mathcal{N}$), the inner product $c_{ji} = \langle T\phi_j, \psi_i \rangle_{\mathcal{N}}$ will have a large (absolute) value whenever $T\phi_j$ and ψ_i correlate, and a small value (in general $\neq 0$) otherwise. The authors showed that an estimate for this diagonal slope can be simply computed as the ratio of areas, $\theta \approx \frac{|\mathcal{M}|}{|\mathcal{N}|}$.

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 $\mathcal{N}$ shown in the inset, and assume a vertex ordering where the points contained in the red region $\mathcal{M}$ appear before those of the blue region $\bar{\mathcal{M}}$. Then, the full Laplacian $\mathbf{L}_{\mathcal{N}}$ will assume the structure

$$\mathbf{L}_{\mathcal{N}} = \begin{pmatrix} \mathbf{L}_{\mathcal{M}} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_{\bar{\mathcal{M}}} \end{pmatrix} + \begin{pmatrix} \mathbf{P}_{\mathcal{M}} & \mathbf{E} \\ \mathbf{E}^\top & \mathbf{P}_{\bar{\mathcal{M}}} \end{pmatrix}, \tag{1.10}$$

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 at the interface between the boundaries $\partial\mathcal{M}$ to $\partial\bar{\mathcal{M}}$.

If the perturbation matrix is identically zero, then (1.10) is exactly block-diagonal; this describes the case in which $\mathcal{M}$ and $\bar{\mathcal{M}}$ are disjoint parts, and the eigenpairs of $\mathbf{L}_{\mathcal{N}}$ are an interleaved sequence of those of the two blocks. The key result shown in [46] is that this interleaving property still holds even when considering the full matrix $\mathbf{L}_{\mathcal{N}}$ as given in (1.10): Its eigenpairs consist of those of the blocks $\mathbf{L}_{\mathcal{M}}$, $\mathbf{L}_{\bar{\mathcal{M}}}$, up to some bounded

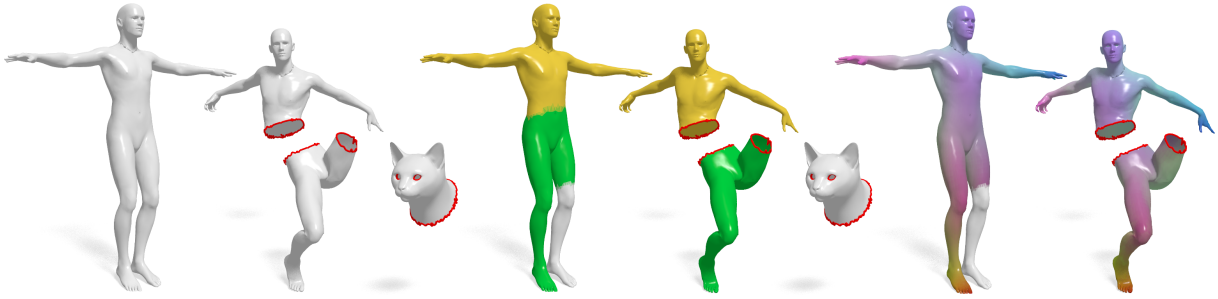

Figure 1.2: Example of the non-rigid puzzle problem considered in Section 1.4: given a model human shape (leftmost, first column) and three query shapes (two deformed parts of the human and one unrelated 'extra' shape of a cat head), the goal is to find a segmentation of the model shape (second column, 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).

perturbation that depends on the length and position of the boundary $\partial\mathcal{M}$. This provides a motivation as to why one observes large correlation $\langle T\phi_j, \psi_i\rangle_{\mathcal{N}}$ with $i \geq j$ in the partial case.

The problem considered in [46] has the form

$$\min_{\mathbf{C},v} \|\mathbf{CA} - \mathbf{B}(v)\| + \rho_{\text{corr}}(\mathbf{C}) + \rho_{\text{part}}(v)\,, \tag{1.11}$$

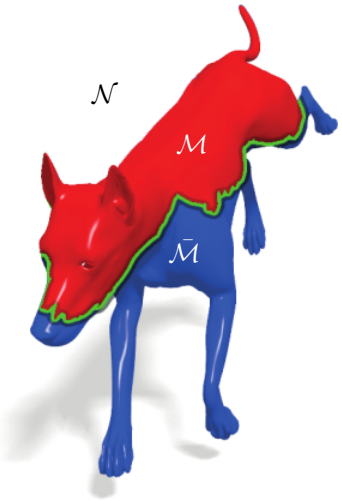

where $v : \mathcal{N} \to [0,1]$ is a (soft) indicator function for the unknown sub-region $\mathcal{N}' \subseteq \mathcal{N}$, and $\mathbf{B}(v) = (\langle \psi_i, v \cdot g_j\rangle_{\mathcal{N}}) = (\langle \psi_i, g_j\rangle_{\mathcal{N}'})$ is the matrix of coefficients for the functions $\{g_i\}$ restricted to the area indicated by v.

The penalties $\rho_{\text{corr}}(\mathbf{C})$ and $\rho_{\text{part}}(v)$ act as regularizers on correspondence and part respectively. The former includes, among several others, a regularization term promoting a slanted diagonal structure on $\mathbf{C}$ with diagonal slope θ, precomputed as the area ratio as discussed above. This way, problem (1.11) incorporates the prior knowledge on the particular structure observed on $\mathbf{C}$ in partial correspondence problems. The $\rho_{\text{part}}(v)$ term favors fewer large contiguous regions over several small fragmented segments, thus imposing a prior on the type of partiality (we refer to [46] for the technical details). Note that function v is defined over the vertices of $\mathcal{N}$, hence scaling linearly with shape size. Problem (1.11) is optimized alternatingly over the Fourier and spatial domains in order to solve for correspondence and part respectively.

1.4 Multi-part partial functional maps

Having set the ground for partial shape matching between a single pair of shapes, in this Section we will describe its extension to multiple parts. Specifically, we are interested in intrinsic, non-rigid, partial, multiple shape correspondence. We shall refer to this setting as *non-rigid puzzles* (see Figure 1.2). We assume to be given a model shape and multiple query shapes, assumed to be parts of (not necessarily 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.

Recently, Litany *et al.* [34] presented a framework for solving 3D non-rigid puzzle problems. They 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. Their method can be considered an extension of [46] for the multiple part setting on one hand, and a non-rigid generalization of the rigid puzzles problem treated in [32] on the other.

In what follows, we will formulate the non-rigid puzzle problem and describe the proposed approach. We will then provide implementation details, followed by experimental results exemplifying how the method copes with some challenging examples.

1.4.1 Non-rigid puzzles

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 $\tau_i : M_i \to N_i$. 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 1.9. By *incompleteness* we mean that the M_i's do not cover $\mathcal{M}$, i.e., there is a *missing part*

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

M_0 can be seen as clutter from the parts perspective. Figure 1.3 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 $\mathbf{\Phi}$ of $\mathcal{M}$ (restricted to M_i) and the Laplacian eigenbasis $\mathbf{\Psi}_i$ of $\mathcal{N}_i$ (restricted to N_i). We further assume to be given as the 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. Since the availability of known corresponding functions is rather a restrictive assumption, in practice it is avoided by 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, descriptors computed on the corresponding $\mathcal{N}_i$'s are used. A robust data fitting term accounts for descriptor mismatches.

With these premises, we formulate the simultaneous segmentation and correspondence

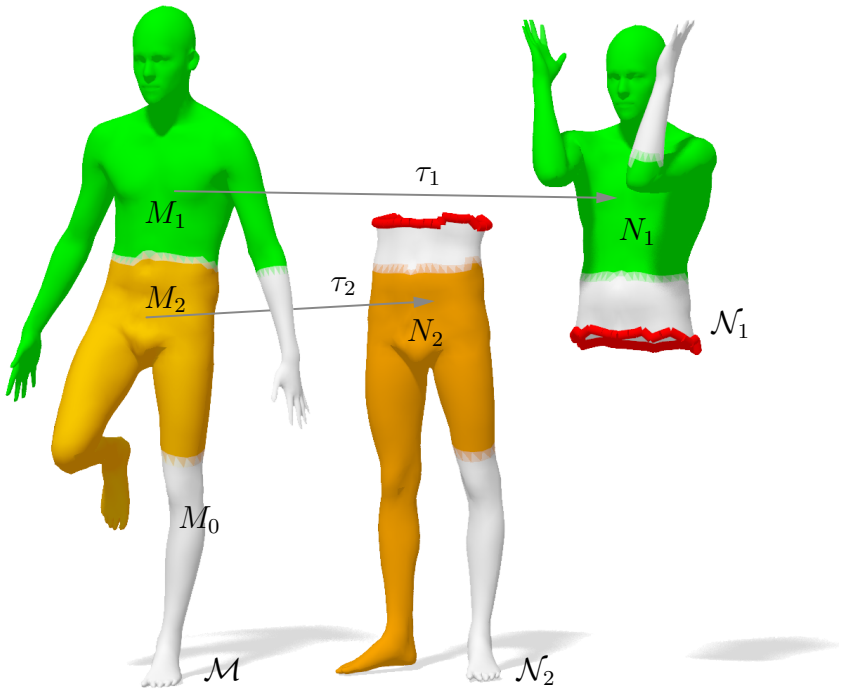

Figure 1.3: The notation we follow in the non-rigid puzzles setting.

as the following optimization problem:

$$\begin{aligned}\min_{\mathbf{C}_i, M_i \subseteq \mathcal{M}, N_i \subseteq \mathcal{N}_i} \quad & \sum_{i=1}^{p} \|\mathbf{C}_i \mathbf{\Psi}_i(N_i)^{\mathrm{T}} \mathbf{G}_i - \mathbf{\Phi}(M_i)^{\mathrm{T}} \mathbf{F}_i\|_{2,1} \\ & + \lambda_{\mathcal{M}} \sum_{i=0}^{p} R_{\mathrm{part}}(M_i) + \lambda_{\mathcal{N}} \sum_{i=1}^{p} R_{\mathrm{part}}(N_i) \\ & + \lambda_{\mathrm{corr}} \sum_{i=1}^{p} R_{\mathrm{corr}}(\mathbf{C}_i) \\ \text{s.t.} \quad & M_i \cap M_j = \emptyset \quad \forall i \neq j \\ & M_0 \cup M_1 \cup \cdots = \mathcal{M} \\ & |M_i| = |N_i| \geq \alpha |\mathcal{N}_i|,\end{aligned} \tag{1.12}$$

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

The first term in (1.12) 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 $R_{\mathrm{part}}(M) = |\partial M|$ 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 segmentation 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 (1.12) is intractable in its combinatorial formulation, the authors proposed a relaxation of the parts to continuous membership functions $u_i : \mathcal{M} \to [0,1]$ to encode the M_i's, and $v_i : \mathcal{N}_i \to [0,1]$ 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

$$\begin{aligned}
\min_{\mathbf{C}_i,\mathbf{u}_i,\mathbf{v}_i} \; & \sum_{i=1}^{p} \|\mathbf{C}_i\mathbf{A}_i(\eta(\mathbf{u}_i)) - \mathbf{B}(\eta(\mathbf{v}_i))\|_{2,1} + \lambda_{\mathcal{M}} \sum_{i=0}^{p} R_{\text{part}}(\eta(\mathbf{u}_i)) \\
& + \lambda_{\mathcal{N}} \sum_{i=1}^{p} R_{\text{part}}(\eta(\mathbf{v}_i)) + \lambda_{\text{corr}} \sum_{i=1}^{p} R_{\text{corr}}(\mathbf{C}_i) \\
\text{s.t.} \quad & \sum_{i=0}^{p} \eta(\mathbf{u}_i) = 1 \\
& \mathbf{a}_{\mathcal{M}}^{\mathrm{T}}\mathbf{u}_i = \mathbf{a}_{\mathcal{N}_i}^{\mathrm{T}}\mathbf{v}_i \geq \alpha \mathbf{a}_{\mathcal{N}_i}^{\mathrm{T}}\mathbf{1}
\end{aligned} \tag{1.13}$$

Here, $\mathbf{a}$ denote the vectors of discrete area elements on the corresponding shapes, and $\eta(t) = \frac{1}{2}\tanh\left(6(t-\frac{1}{2})\right) + \frac{1}{2}$ is an element-wise non-linear transformation used to restrict the indicators at each vertex to the range $[0,1]$. The matrices $\mathbf{A}_i(\eta(\mathbf{u}_i)) = \mathbf{\Phi}^{\mathrm{T}}\mathrm{diag}(\mathbf{u}_i)\mathbf{F}_i$ and $\mathbf{B}_i(\eta(\mathbf{v}_i)) = \mathbf{\Psi}_i^{\mathrm{T}}\mathrm{diag}(\mathbf{v}_i)\mathbf{G}_i$ denote the representation coefficients of the input corresponding functions restricted to their respective parts.

As the regularization term of the segments the authors make use of a discretized version of the intrinsic Mumford-Shah functional introduced in [9]

$$R_{\text{part}}(\mathbf{u}) = \int_{\mathcal{M}} \xi(\mathbf{u}) \|\nabla_{\mathcal{M}} \mathbf{u}_i\| da \approx \mathbf{a}_{\mathcal{M}}^{\mathrm{T}} \mathbf{g}_i \tag{1.14}$$

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

Following [46] for the regularization of the functional maps $\mathbf{C}_i$,

$$\begin{aligned}
R_{\text{corr}}(\mathbf{C}) &= \|\mathbf{C} \odot \mathbf{W}\|_{\mathbf{F}}^2 + \lambda_1 \sum_{i \neq j} (\mathbf{C}^{\top}\mathbf{C})_{ij}^2 \\
&+ \lambda_2 \sum_{i} ((\mathbf{C}^{\top}\mathbf{C})_{ii} - d_i)^2.
\end{aligned} \tag{1.15}$$

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

$$w_{ij} = e^{-\sigma\sqrt{i^2+j^2}} \| \frac{\mathbf{n}}{\|\mathbf{n}\|} \times ((i,j)^{\top} - \mathbf{p}) \|.$$

The slanted diagonal of $\mathbf{W}$ is a line segment $\delta(t) = \mathbf{p} + t\frac{\mathbf{n}}{\|\mathbf{n}\|}$ with $t \in \mathbb{R}$, where $\mathbf{p} = (1,1)^{\top}$ is the matrix origin, and $\mathbf{n} = (1, |\mathcal{N}_i|/|\mathcal{M}|)^{\top}$ is the line direction with slope $|\mathcal{N}_i|/|\mathcal{M}|$.

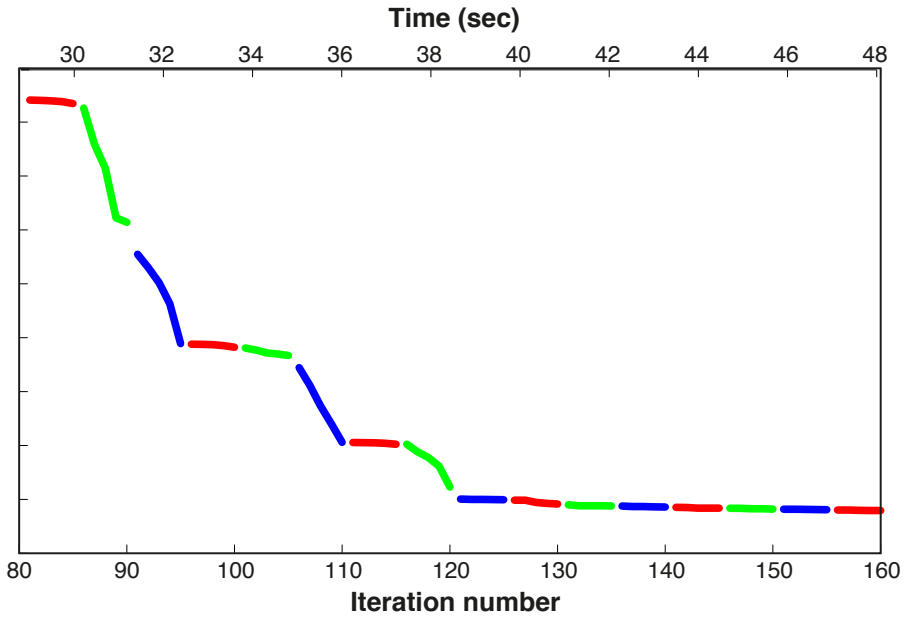

Figure 1.4: An example showing the decrease in cost during the alternating minimization of the different sets of variables, $\mathbf{C}_i$ (red), u_i (green) and v_i (blue).

The second factor in w_{ij} is the distance from the slanted diagonal δ, and $\sigma > 0$ regulates the spread around δ. In all experiments shown, σ was set to a fixed value of 0.03. The second term in $R_{\text{corr}}(\mathbf{C})$ promotes orthogonality of $\mathbf{C}$, while the third term, setting the first few d_i to 1 and the rest to 0, regularizes its rank.

1.4.2 Implementation

The resulting optimization problem is solved by means of a threefold alternating minimization. To this end, the masks $\mathbf{W}_i$ are used to initialize the matrices $\mathbf{C}_i$ by applying the transformation $\mathbf{C}_{i_{m,n}} = 1 - \frac{\mathbf{W}_{i_{m,n}}}{max(W_i)}$. Then a minimization over the partial functional maps $\mathbf{C}_i$, the model indicator functions u_i and the parts indicator functions v_i is performed in a cyclic manner, keeping the other parameters fixed. Although this algorithm is not guaranteed to converge, in practice the authors reported strictly decreasing cost value as the one shown in Figure 1.4. For the different minimization steps a conjugate-gradient solver was used, as supplied by the Manopt toolbox [6]. 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 mapping $\mathbf{C}_i$, a k-dimensional ICP was added [46]. As noted by the authors, this step helps especially when the descriptors are performing poorly. The parameters were changed according to required setting of the experiment. For instance, in the non-isometric experiment λ_2 was set to 0 to allow change of areas.

1.4.3 Experimental results

The non-rigid puzzles method (NRP) 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).

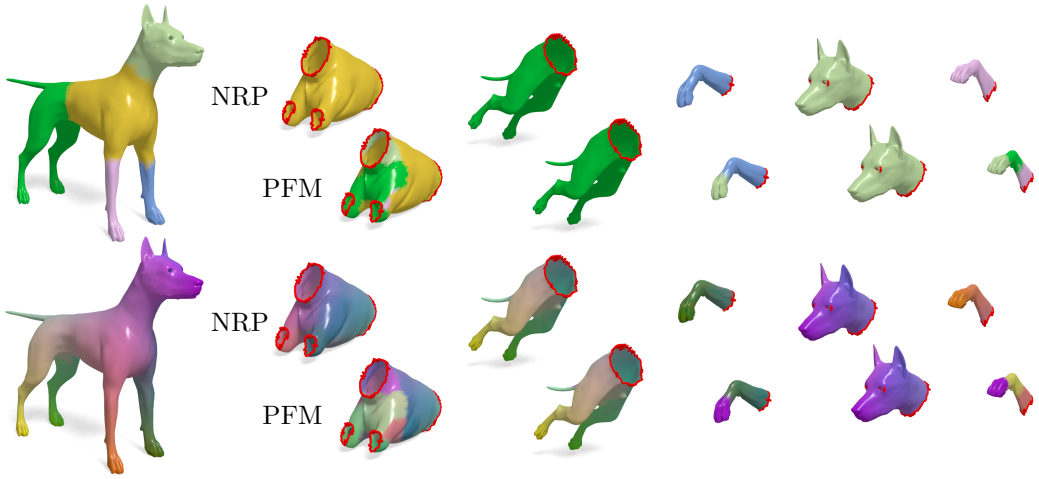

Figure 1.5: Comparison between partial functional maps [46] and NRP in a perfect puzzle setting. For each method we show the membership functions of each part with respect to the model (first two rows), and the color-coded correspondence between parts and model (last two rows). For PFM, the fact that 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 NRP as all parts are matched jointly to the template, yielding a regularizing effect on the correspondence.

Data. Experiments include both synthetic and real data. The synthetic dataset is made up of shapes from the TOSCA [8] and FAUST [5] benchmarks. In order to avoid compatible meshings and make the dataset more realistic, each TOSCA model is independently remeshed to ∼10K vertices by iterative pair contractions [20]. All FAUST templates are kept at their original resolution (∼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 [24].

The shapes from these datasets are decomposed into a controllable amount of parts by Voronoi decomposition and consensus segmentation [45]; 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 descriptors for the data term in (1.12) a 352-dimensional SHOT signatures [54] were used. 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 [3, 53]. Note that we compute dense descriptors for all shape points, including those lying along the boundaries.

Perfect puzzle. Figure 1.5 depicts an example of a solution obtained with NRP 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.* [46] 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 1.6.

Overlapping pieces. A more interesting setup is obtained when allowing the different pieces to have non-zero overlap. In Figure 1.8 we show the 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 NRP 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.

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 1.4.1, NRP 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 1.2 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 1.9.

1.4.4 Discussion and conclusions

In this section we described a method for solving 3D non-rigid puzzle problems. Th problem was formulated 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. This framework is quite flexible, and can be easily adapted to deal with missing or overlapping pieces, moderate amounts of clutter, and outliers.

Limitations. One of the main limitations which NRPinherits from the functional maps framework [41], 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. 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. Equally important, and this time inherited from the *partial* functional maps [46], is the use of an indicator function per part, resulting in substantial computational complexity. A possible solution to this limitation is described in the following Section.

1.5 Fully-spectral partial functional maps

The main drawback of partial functional maps [46] and the follow-up works [17, 34] is their explicit model of the part, requiring a somewhat cumbersome solver alternating between optimization in the spatial domain (over the part indicator function) and in the spectral domain (over the correspondence matrix). Furthermore, the complexity of the spatial domain optimization depends on the number of mesh vertices and scales poorly (see Section 1.5.2 for an evaluation). The method presented in this Section builds on a simple observation allowing to formulate the partial functional maps problem *entirely in the spectral domain*. It bears resemblance to joint approximate diagonalization but has a fundamental difference that will be emphasized in the sequel.

Localization. A key feature of partial functional maps lies in their *spatially localized* behavior: Any solution to (1.11) is a map $T : L^2(\mathcal{M}) \to L^2(\mathcal{N})$ that is supported on some region $\mathcal{N}' \subseteq \mathcal{N}$ of the full model, meaning that for all $y \in \mathcal{N} \setminus \mathcal{N}'$ the approximate equality $(Tf)(y) \approx 0$ holds for any $f \in L^2(\mathcal{M})$. This can be easily seen by noting that the image of $\mathbf{A}$ under $\mathbf{C}$ must be localized to the region indicated by v in order for the data term $\|\mathbf{CA} - \mathbf{B}(v)\|$ to reach a minimum; in other words, the functional map $\mathbf{C}$ must localize the correspondence.

This localization property comes at the price of modeling the region $\mathcal{N}' \subseteq \mathcal{N}$ explicitly. In what follows we describe the method proposed by [33] to absorb the spatial mask into a new basis $\{\hat{\psi}_j\}$ for $L^2(\mathcal{N})$; in doing so, one disposes of the explicit part v and obtains a simpler optimization problem, as elucidated in the following.

Assume $\mathbf{C}, v$ are a solution to (1.11), such that $\mathbf{CA} = \mathbf{B}(v)$ holds approximately, and consider two functions $f \in L^2(\mathcal{M}), g \in L^2(\mathcal{N})$ whose spectral representations are columns of $\mathbf{A}$ and $\mathbf{B}$ respectively. In the spatial domain, the equality becomes

$$\sum_{ij}^{k} \langle f, \phi_i \rangle_{\mathcal{M}} c_{ji} \psi_j = \sum_{i=1}^{k} \langle v \cdot g, \psi_i \rangle_{\mathcal{N}} \psi_i \approx v \cdot g \,, \tag{1.16}$$

where the approximation is due to truncation to the first k terms. By defining a new basis $\hat{\psi}_i = \sum_{j=1}^{k} c_{ji} \psi_j$, we get to

$$\sum_{ij}^{k} \langle f, \phi_i \rangle_{\mathcal{M}} \hat{\psi}_j \approx v \cdot g \,, \tag{1.17}$$

in other words, the modified basis $\{\hat{\psi}_j\}$ induces the sought localization. Importantly, in order for (1.17) to hold for general f and g, the new basis functions themselves must be localized, *i.e.*, $\hat{\psi}_i = v \cdot \hat{\psi}_i$ for all i.

Using the fact that orthogonal $\mathbf{C}$ implies orthogonal $\{\hat{\psi}_j\}$, Equation (1.17) can be phrased in the spectral domain as:

$$\mathbf{A} \approx \mathbf{CTB}(v) = \mathbf{CTB} \,; \tag{1.18}$$

in the last equality, the indicator function v was absorbed into the new basis functions $\{\hat{\psi}_j\}$.

Problem. In light of our previous analysis, the following manifold optimization problem is considered:

$$\min_{\mathbf{Q}\in S(k,r)} \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{A}_r - \mathbf{Q}^\top \mathbf{B}\|_{2,1}\,, \tag{1.19}$$

where $S(k,r)$ is the Stiefel manifold of orthogonal $k \times r$ matrices (ortho-projections), and $\mathbf{A}_r = \mathbf{W}_r \mathbf{A}$ with $\mathbf{W}_r = (\mathbf{I}_{r\times r}\ \mathbf{0}_{r\times k-r})$ denotes the $r \times k$ matrix containing the first r rows of $\mathbf{A}$. The value of r is directly related to the rank of the partial functional map $\mathbf{C}$ in (1.11) and can be estimated simply from the area ratio θ, or optimized for explicitly by solving (1.19) for a range of r's. The rank r and the orthogonality of $\mathbf{Q}$ act as partiality priors, since they are related to the underlying map being area-preserving [41, 46].

The optimization problem (1.19) models partial correspondence as the search for a new basis that is localized to a latent part of the full shape. In this view, the matrix $\mathbf{Q}$ is not regarded as a functional map between shapes, but rather as a matrix of transformation coefficients for the basis (the off-diagonal regularity term ensures that the transformation is smooth). This interpretation allow one to tackle part-to-part settings (see Eq. (1.21)) as a simple modification to (1.19). The first r functions $\{\hat{\psi}_1, \dots, \hat{\psi}_r\}$ of the new orthogonal basis $\hat{\psi}_i = \sum_{j=1}^{k} q_{ji}\psi_j$ obtained as the result of such a transformation would be approximately orthogonal to $\{\phi_1, \dots, \phi_r\}$ under the functional correspondence (see Figure 1.1),

$$\langle T\phi_i, \hat{\psi}_j \rangle_{\mathcal{N}} \quad \approx \quad \delta_{ij}; \quad i,j = 1, \dots, r\,. \tag{1.20}$$

It is important to remark that, while the correct partial correspondence is a solution to our problem by Eq. (1.16–1.18), this is not necessarily unique as it directly depends on the input data. Not all such optima are localized to the correct region, and some might even have global support. The choice of the input corresponding functions $\{f_i, g_i\}$ ultimately determines the quality of the localization (see Figure 1.13). In practice, as reported by the authors of [33] it is enough to employ dense descriptor fields that are sufficiently similar on the corresponding regions in order to drive the optimization to the correct solution.

Part-to-part. Let us now assume that only a part $\mathcal{N}'$ of the shape $\mathcal{N}$ matches the corresponding part $\mathcal{M}' \subset \mathcal{M}$ (see Figure 1.13). As observed by Litany *et al.* [34], one still obtains a slanted-diagonal structure of $\mathbf{C}$ with angle $\theta = \frac{|\mathcal{M}|}{|\mathcal{N}|}$, that is, θ depends only on the area ratio of the known full shapes and not that of the unknown parts, $\frac{|\mathcal{M}'|}{|\mathcal{N}'|}$. On the other hand, if $\mathcal{N}'$ were given, only about $\frac{|\mathcal{N}'|}{|\mathcal{M}|}k$ out of k first eigenfunctions of $\Delta_{\mathcal{M}}$ would correspond to the first k eigenfunctions of $\Delta_{\mathcal{N}}$. This means that while the matrix $\mathbf{C}$ in the partial functional correspondence problem (1.11) will have the same slanted diagonal structure regardless of the size of the corresponding parts $\mathcal{M}'$ and $\mathcal{N}'$, the actual fraction of non-zero entries on the slanted diagonal will be about $\min\left\{\frac{|\mathcal{N}'|}{|\mathcal{M}|}, \frac{|\mathcal{M}'|}{|\mathcal{N}|}\right\}$ or, assuming approximately isometric parts, $\min\left\{\frac{|\mathcal{M}'|}{|\mathcal{M}|}, \frac{|\mathcal{N}'|}{|\mathcal{N}|}\right\}$. Moreover, the exact indices of these corresponding functions cannot be predicted *a priori*.

Since the first r eigenfunctions of $\Delta_{\mathcal{M}}$ typically contain only a subset of all the corresponding eigenfunctions, in order to satisfy (1.19) we have to modify the coefficients $\mathbf{A}$ as well. This leads to

$$\min_{(\mathbf{P},\mathbf{Q})\in S^2(k,r)} \operatorname{off}(\mathbf{P}^\top \mathbf{\Lambda}_{\mathcal{M}} \mathbf{P}) + \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}\,, \tag{1.21}$$

where optimization is now performed on the *product of Stiefel manifolds*. Figure 1.13 illustrates the localization behavior of the new bases under different inputs $\mathbf{A}, \mathbf{B}$, and as a function of r. As we will also show in the experimental section, it is sufficient to use robust enough descriptor fields in order to get a good localization to the latent corresponding region. Further note how the choice of r also affects map locality.

Comparison to joint diagonalization. Problems (1.19) and (1.21) can be viewed as variants of joint approximate diagonalization problems (1.9) and (1.8), respectively with the ℓ_2 data fitting term replaced by the more robust $\ell_{2,1}$ counterpart as was previously suggested in [43] and [27]. Despite this resemblance, the crucial difference lies in the fact that in the former problems $k \times r$ ortho-projections are used in place of full-rank $k \times k$ orthogonal matrices.

The data term of problem (1.21) can be rewritten using full-rank $k \times k$ orthogonal matrices $\mathbf{P}, \mathbf{Q}$ as $\|\mathbf{W}_r(\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B})\|_{2,1}$ and can be interpreted as the fitting term of (1.8) with a modified metric. The effect of using the mask $\mathbf{W}_r$ is visualized in Figure 1.15.

Theoretical and experimental justification provided in this paper suggests that, surprisingly, such an apparently simple modification of the problem is sufficient to handle a wide range of settings involving partiality, clutter, and topological noise, as well as lack of isometry.

Geometric interpretation. The joint approximate diagonalization process (1.8) can be interpreted as a *rigid alignment* of the k-dimensional spectral embeddings $\{\phi_i\}_{i=1}^k$ and $\{\psi_i\}_{i=1}^k$ of two shapes, where the orthogonal matrices $\mathbf{P}, \mathbf{Q}$ rotate/reflect the eigenfunctions such that the resulting bases (1.7) are aligned. Similarly, the approach presented here (1.21) can be interpreted as a *non-rigid alignment* in the r-dimensional eigenspace. The new bases $\{\hat{\phi}_i\}_{i=1}^r$ and $\{\hat{\psi}_i\}_{i=1}^r$ are constructed as linear combinations of k eigenvectors; if $k \gg r$, one can produce almost arbitrary sets of r aligned orthogonal basis functions. The off term in problem (1.21) acts as a regularization ensuring that the functions are smooth. The combined effect of the data and regularization terms is that of a non-rigid alignment (see Figures 1.11, 1.12 and 1.15). Note that while some spectral approaches [23,36,47] seek for a correspondence by non-rigid ICP in the spectral domain, none of these successfully tackle the case of missing geometry and topological noise.

Comparison to partial functional maps. As discussed earlier, any solution to the functional correspondence problem (1.11) is also a solution to FSPM (1.19). A key difference lies in the direction of the map: If we regard matrix $\mathbf{Q}$ as the spectral representation of a functional map, the data term of FSPM evaluates its pre-image in $\mathcal{M}$ (the partial shape), while that of (1.11) looks at the image on the full shape $\mathcal{N}$, thus requiring an explicit modeling of the part.

Further, as described in [46], the regularizer $\rho_{\text{corr}}(\mathbf{C})$ in (1.11) includes a penalty term promoting $\mathbf{C}^\top \mathbf{C} \approx \begin{pmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. Here, this area-preservation requirement is instead phrased as a hard constraint, where optimization is done over the Stiefel manifold $S(k, r)$, such that $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_r$. Overall, the optimization problem (1.19) is less engineered than (1.11), has less parameters, and is simpler to optimize. See Figures 1.14-1.16 for further comparisons.

1.5.1 Implementation

Problem (1.19) is a manifold-constrained and non-smooth (due to the $\ell_{2,1}$ norm). It is solved using the MADMM scheme of Kovnatsky *et al.* [27].

The authors have implemented the solver in Matlab using `manopt` [6], a framework for optimization over manifolds. Laplacians were discretized using the classical cotangent scheme [37, 42]. Code is publically availiable at `https://github.com/orlitany/FSPM`note that although this step and the subsequent eigen-decomposition clearly depend on the number of vertices, they are carried out only once for each shape and thus count as an off-line cost. Although the manifold constraints render the problem non-convex and MADMM gives no global optimality guarantees, in practice the authors have observed a stable behavior with a strictly decreasing cost value and fast convergence (an empirical evaluation is provided in Section 1.5.2).

Initialization. The orthogonal matrices in problems (1.19) and (1.21) are initialized as $k \times r$ random matrices with $k = 50$ and r estimated via the area ratio. Since the availability of known corresponding functions for the data term is a restrictive assumption, in practice $\{f_i, g_i\}_{i=1}^{q}$ is replaced with dense descriptor fields calculated on $\mathcal{M}$ and $\mathcal{N}$, where q is the number of dimensions of the descriptor. In all experiments the 352-dimensional SHOT [54] were used with default parameters.

Point-wise map conversion and refinement. After convergence, the point-wise correspondence are recovered by a nearest-neighbor search in the k-dimensional spectral domain [41] (*e.g.*, in the example of Figure 1.15, each blue point is matched to the closest red point). The solution is further refined by selecting 10% of the matches using farthest point sampling, and using them to construct new corresponding functions $\{f_i, g_i\}$ as sparse (yet well spread) localized smooth delta functions. The new data term replaces the initial one, which was based solely on descriptors. The value of μ is adjusted accordingly to keep a similar weight between the new data term and the regularizer. This process is repeated 5 times. For a fair comparison, the same refinement procedure was applied to partial functional maps [46] and joint diagonalization [29]. Note that while more sophisticated recovery methods exist [47], these work under the assumption of no partiality.

1.5.2 Experimental results

The proposed method, dubbed FSPM, was evaluated extensively in a variety of settings. It was executed on an Intel i7-4710MQ 2.50GHz CPU with 8 logical cores.

Evaluation. Correspondence quality is quantitatively evaluated according to the Princeton benchmark protocol [25]. Assume that a correspondence algorithm produces a match $(x, y) \in \mathcal{M} \times \mathcal{N}$, whereas the ground-truth correspondence is (x, y^*). Then, the inaccuracy of the correspondence is measured as

$$\epsilon(x) = \frac{d_{\mathcal{N}}(y, y^*)}{|\mathcal{N}|^{1/2}}\,, \tag{1.22}$$

where $d_{\mathcal{N}}$ is the geodesic distance on $\mathcal{N}$. Cumulative curves are plotted, showing the percent of matches which have error smaller than a variable threshold. Symmetric solutions are given no penalty.

Topological changes A full quantitative evaluation was performed on the recent SHREC'16 Topology benchmark [30] (*low resolution* setting, ∼10K vertices per shape). The dataset consists of 90 matching problems between human shapes undergoing topological changes of various intensity (some examples are shown in Figure 1.23). The methods appearing in the original benchmark are random forests (RF) [48], Green's embedding (GE) [30], and isometric embedding (EM) [52]. As reported in Figure 1.17, previous approaches demonstrated poor performance due to the challenging setting. A comparison with partial functional maps (PFM) [46] was additionally included, as well as with recent convex optimization (CO) method of Chen and Koltun [14], which performs an explicit modeling of topological artifacts but did not previously appear in the benchmark. For a fair comparison, the extrinsic regularization term of [14] was disabled since it relies on the shapes being approximately aligned in $\mathbb{R}^3$.

Part-to-full. A quantitative evaluation of FSPM in the partial matching scenario was done on the challenging SHREC'16 Partial Correspondence benchmark [18]. The dataset is composed of 400 partial shapes (from a few hundred to ∼9K vertices each) belonging to 8 different classes (humans and animals), undergoing nearly-isometric deformations in addition to having missing parts of various forms and sizes. Each class comes with a "null" shape in a standard pose which is used as the full template to which partial shapes are to be matched. This results in 400 matching problems in total. The dataset is split into two subsets, namely *cuts* (removal of a few large parts) and *holes* (removal of many small parts).

The results are reported in Figures 1.18 and 1.19, and qualitatively in Figure 1.24. As a comparison, the authors included results for partial functional maps (PFM) [46], random forests (RF) [48], scale-invariant isometric matching (IM) [51], game-theoretic matching (GT) [44], and elastic net matching (EN) [49], as these methods appeared in the original benchmark. Additionally, joint diagonalization (JAD) [29] was included in the comparison.

It can be seen from the plots that FSPM method has a ∼10% improvement on PFM, the closest competitor, in both datasets. Given the purely spectral nature of FSPM and its considerably simpler formulation, as opposed to the cumbersome optimization in the spatial domain performed by PFM, we find these results quite remarkable (a runtime comparison of the two methods will be presented in Section 1.5.2) The poor performance of JAD puts in evidence the importance of correctly estimating the rank of the new basis.

Scanned data. Qualitative experiments were carried out on the FAUST dataset [5], which contains real human shapes acquired with a 3D scanning device. By nature of the acquisition process, these shapes are affected by topological artifacts as well as missing parts due to self-occlusions, resulting in a challenging testbed for shape matching. The results are shown in Figure 1.22.

Runtime. In Figure 1.21 runtime comparison with PFM are reported at increasing number of vertices. Since FSPM technique operates exclusively in the spectral domain, the computational cost of each iteration only depends on the prescribed basis dimension k, hence it is constant w.r.t. shape size (see Equation (1.19)). In contrast, due to the alternating optimization over the spectral and spatial domains, the runtime complexity of PFM grows linearly with shape size (Equation (1.11)). The average runtime on the SHREC'16 benchmarks was ∼220 sec. for FSPM method and ∼1240 sec. for PFM.

1.5.3 Discussion and conclusions

In this section we have discussed FSPM, a recent method for partial dense intrinsic correspondence between deformable shapes. Contrarily to previous approaches, FSPM is generic in that it allows to tackle topological noise, strong partiality, and non-isometric deformations within the same framework, making it amenable for application in practical settings involving real data acquisition. A remarkable feature of FSPM lies in its purely spectral nature, allowing to perform all calculations (except for the initial calculation of the first k Laplacian eigenfunctions) with constant complexity independent of the shape size. FSPM improves the state of the art for shape correspondence on three recent benchmarks, where it is faster than the closest competitor by one order of magnitude, and performs demonstrably well on real data.

Limitations. Examples of failure cases are shown in Figure 1.20. The main limitation of FSPM lies in its reliance on good local features to drive the matching process. Correspondence quality is directly affected by the robustness of the chosen descriptor fields to the artifacts that one may encounter in practice, and designing a local descriptor that is robust to deformations, topological noise and missing geometry is indeed an open challenge tackled by few. Second, our approach shares with other intrinsic methods its invariance to intrinsic symmetries, resulting in reflected solutions that may be undesirable in certain applications. Operating again at the feature level by incorporating some notion of symmetry-awareness (hence an extrinsic quantity) in the local descriptor may be a possible and promising direction to pursue.

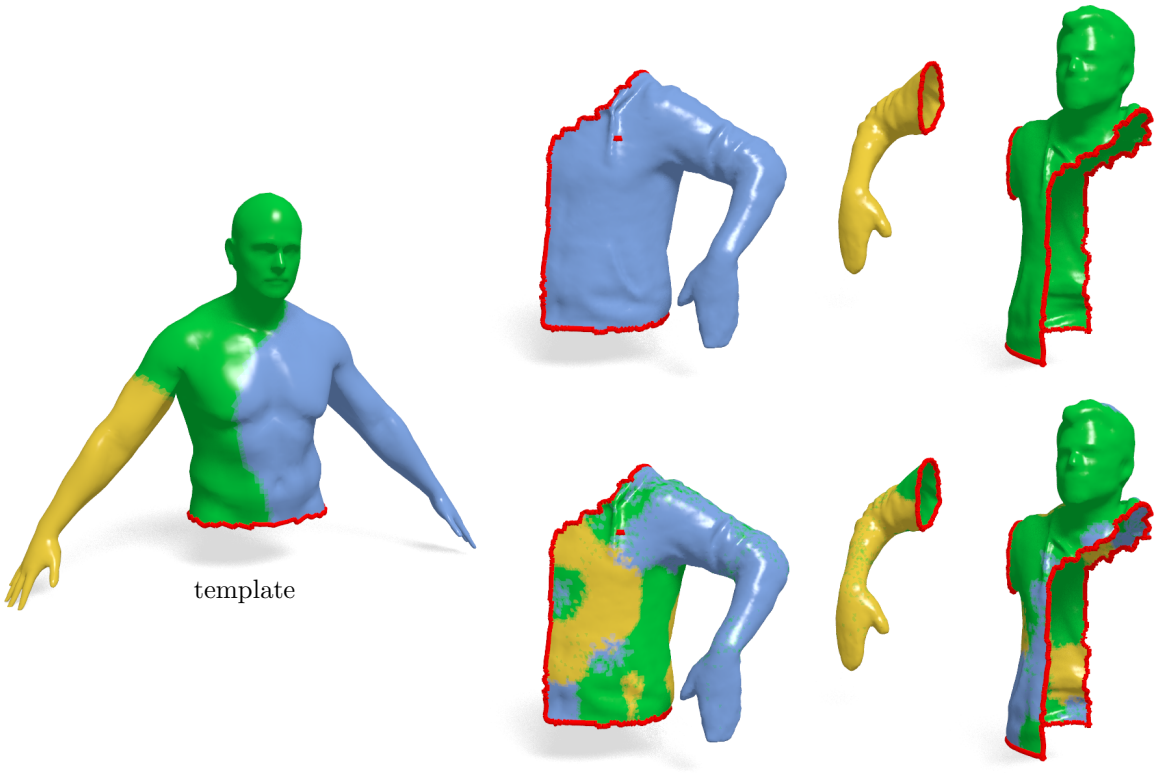

Figure 1.6: Comparison between NRP (top row) and PFM (bottom row) on real data. The parts shown on the right were acquired with a 3D scanner.

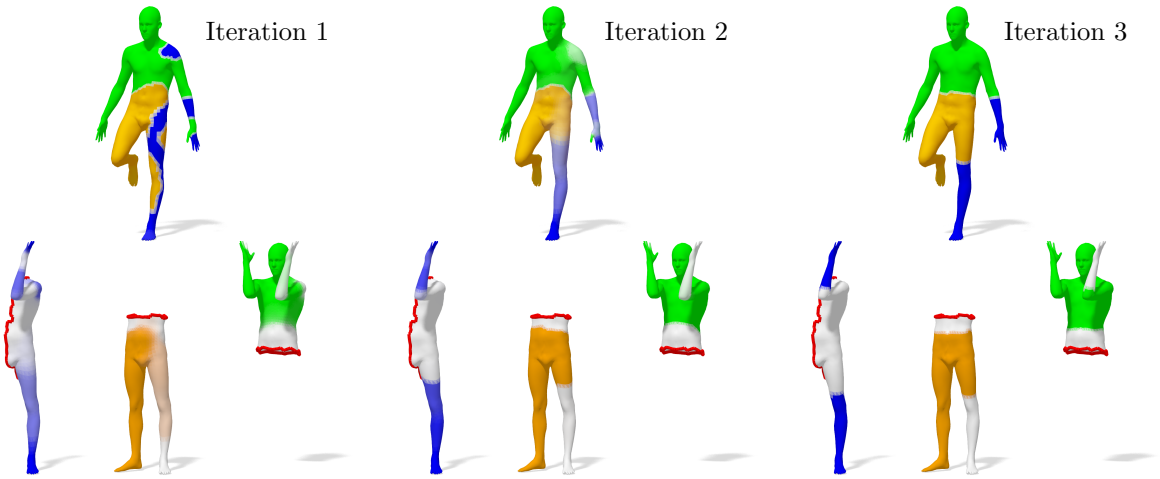

Figure 1.7: An example of NRP matching pipeline when dealing with overlapping parts. The optimization process alternates between the membership functions on the model (top row) and those on the parts (bottom row). At each alternating step, the membership functions are optimized jointly on the respective parts. Note that in this particular example there is more than one possible solution, *e.g.*, the blue part could be completely excluded from the solution.

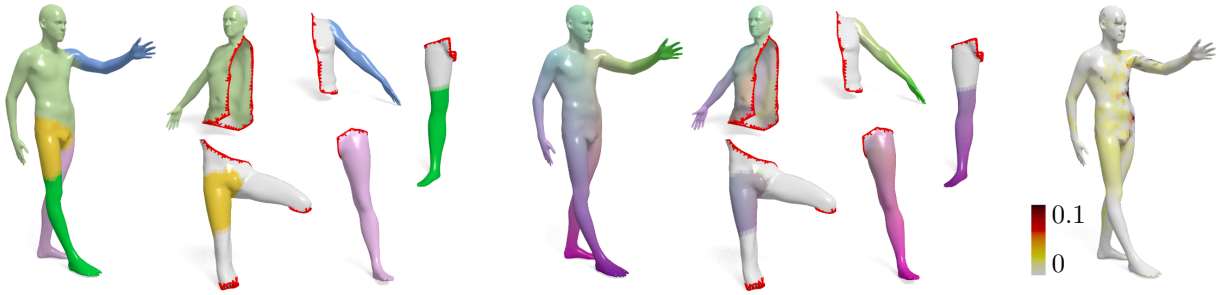

Figure 1.8: Non-rigid puzzle alignment between overlapping parts. Shown are the final segmentation obtained by NRP (left), the dense matchings between the parts and the model (middle), and the normalized geodesic error (shown as a heatmap) to the ground-truth correspondence (right). With the exception 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, NRP 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.

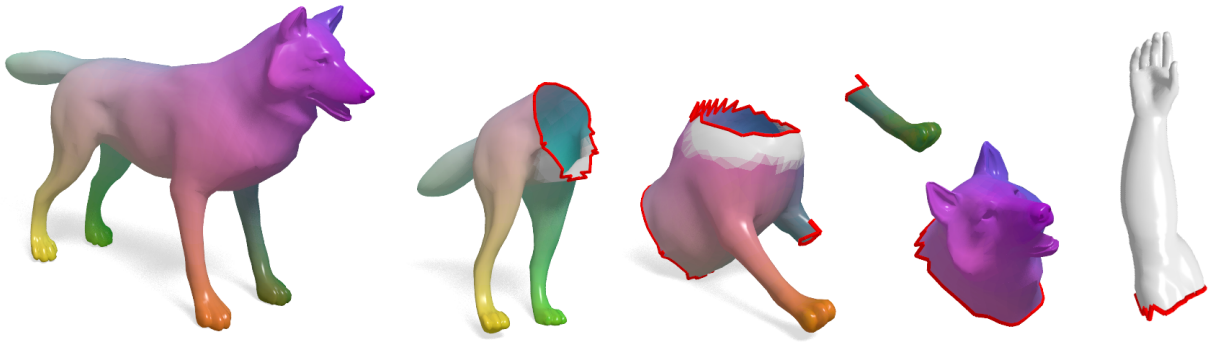

Figure 1.9: In this example, an additional outlier piece (the human arm) is included in the input set. NRP treats extra pieces as clutter; the arm is automatically selected by the matching process, and completely excluded from the final solution. Note how the presence of the extraneous part did not affect the quality of the correspondence.

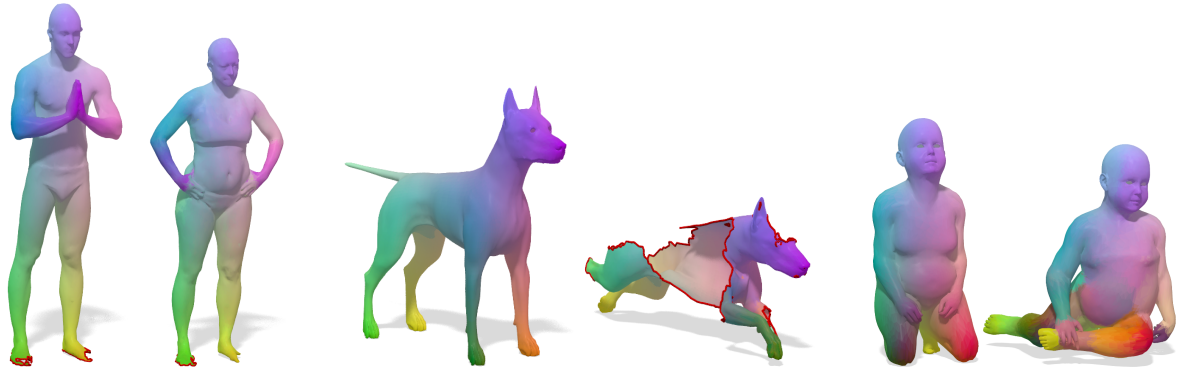

Figure 1.10: Examples of dense correspondence computed with FSPM on real 3D scans (left pair, the areas of contact are glued together), missing parts (middle) and strong topological artifacts (right, touching parts are glued together). Corresponding points are encoded with the same color.

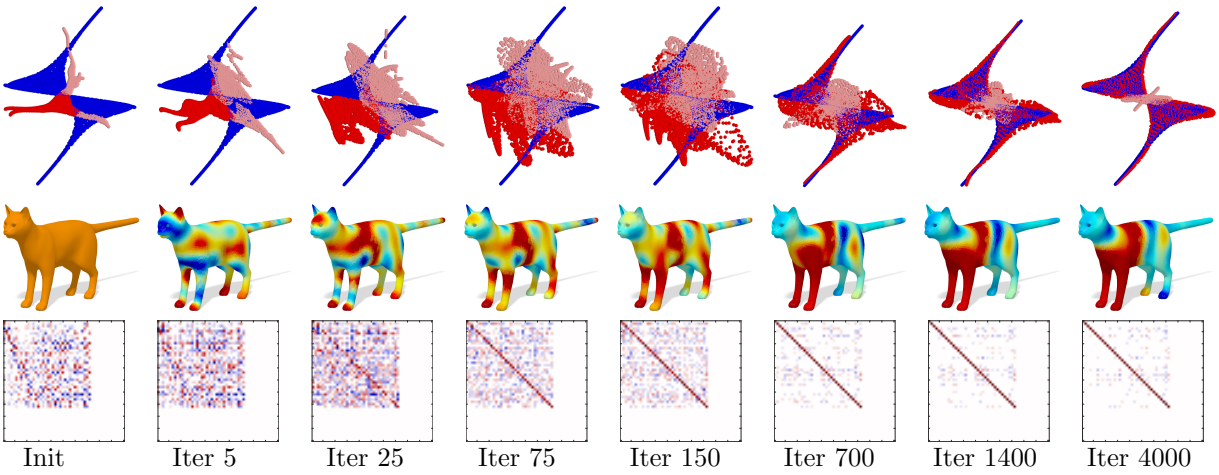

Figure 1.11: The optimization process admits an interpretation as a non-rigid alignment of the spectral embeddings of the input shapes (in this example, the cat meshes of Figure 1.15). *Top*: The spectral embeddings (in 2D for simplicity) at different time steps. The localization effect is manifested in the "extra" unmatchable part (pink point cloud) shrinking towards zero in the spectral domain as the optimization converges to a correct partial correspondence. In the last iterations, the two matching point clouds (blue and red) are almost perfectly aligned. *Bottom*: Correspondence matrices in the new basis, computed as the matrix $\mathbf{C}(t)$ minimizing $\mathbf{C}(t)\mathbf{A} = \mathbf{Q}(t)^\top\mathbf{B}$ in the least squares sense. Note that this correspondence matrix is never actually used in the matching pipeline, and is being included here for illustration purposes. *Middle*: A segmentation function indicating the region on the full shape that is put into correspondence with the partial shape (cold and hot colors represent small and large values respectively). The function is simply taken to be the image of the constant function via $\mathbf{C}(t)$.

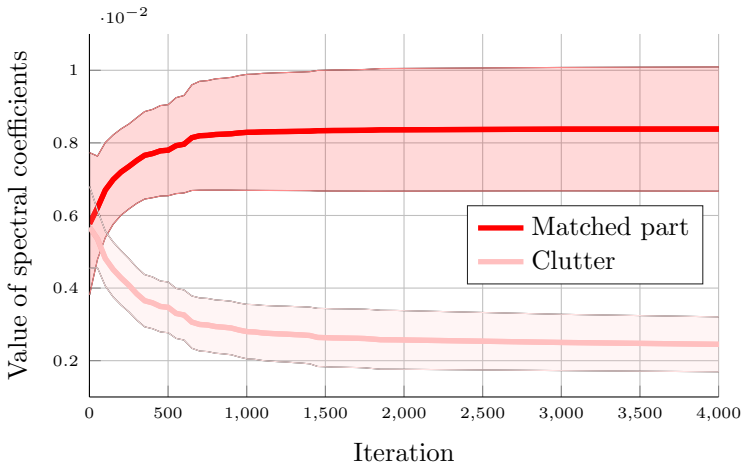

Figure 1.12: This plot shows the evolution of the spectral coefficients of points belonging to the matchable (in red) and unmatchable (in pink) parts for the example shown in Figure 1.11. Observe the formation of two distinct groups of values from the very first iterations, with the values of the unmatched part tending towards zero.

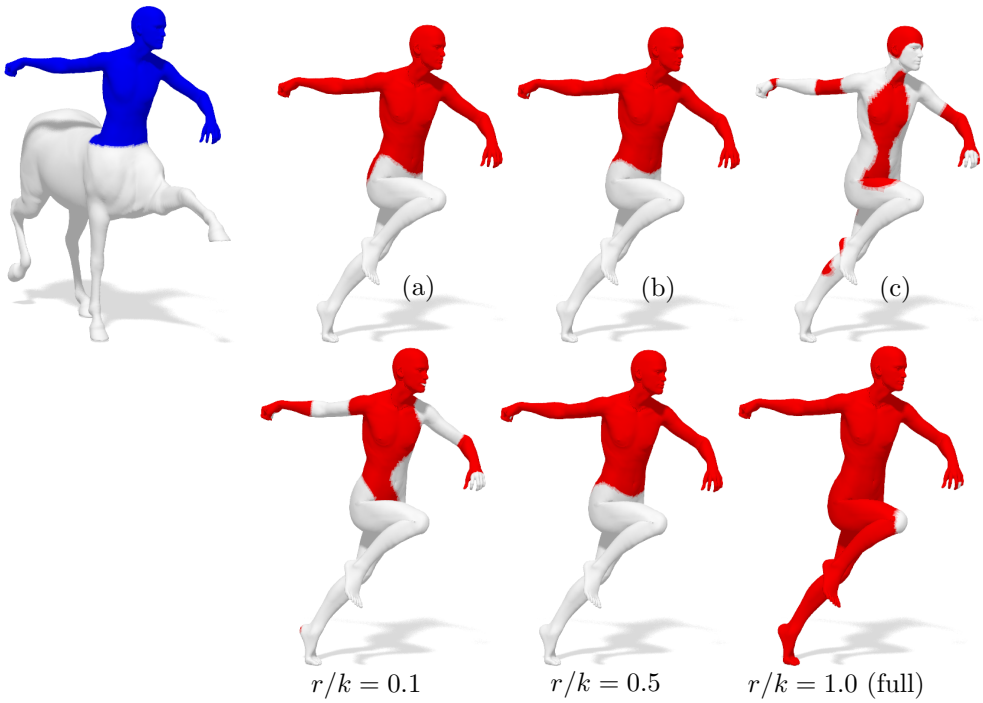

Figure 1.13: Effect of changing the data term in (1.21) by using *(a)* descriptor fields localized to the correct region of the human shape (*i.e.*, the region corresponding to the human part of the centaur); *(b)* descriptor fields supported on the entire shape, but similar only on the correct region; and *(c)* noisy descriptors with similar values outside the correct region. The segmentation of the human shape is visualized after optimization of (1.21). In the bottom row, a localization is shown as a function of the rank r.

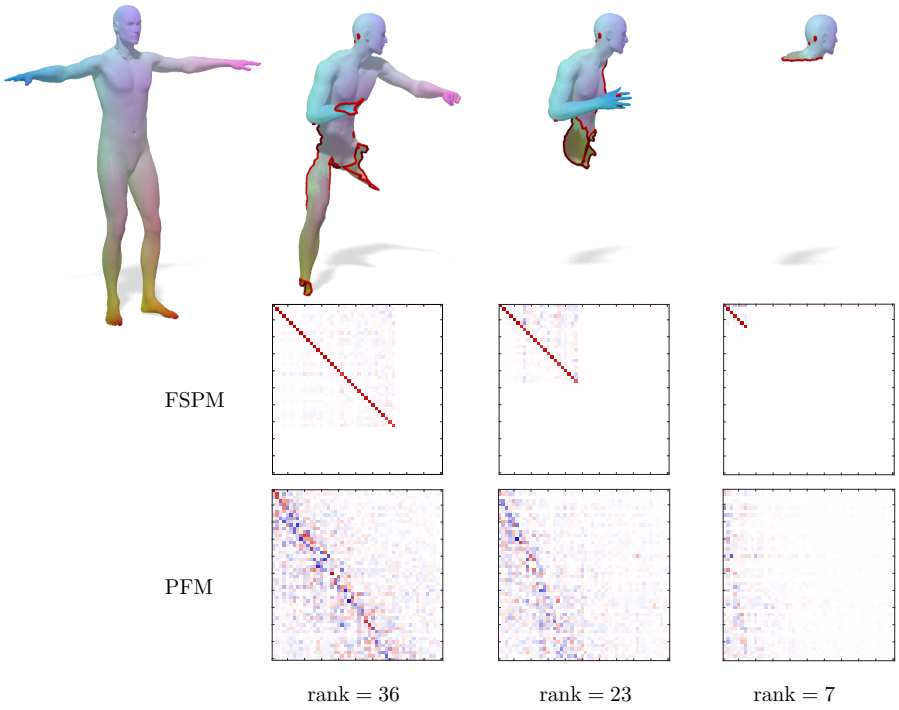

Figure 1.14: Correspondence matrices at increasing partiality. Shown are solutions obtained by FSPM in the new basis (middle row) and by PFM in the standard Laplacian eigenbasis (bottom row). Observe how FSPM representation remains crisp even at extreme levels of partiality (rightmost column).

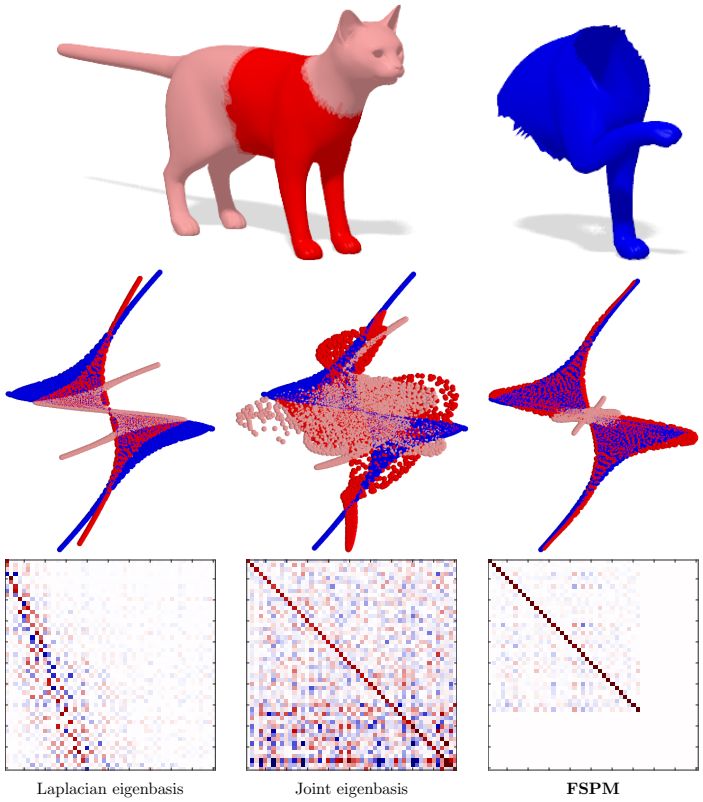

Figure 1.15: *Top*: Input shapes (dark red denotes the part corresponding to the partial blue shape). *Middle*: Spectral embeddings (shown are the first two eigenfunctions) using the standard Laplacian eigenbasis (left), in the basis obtained by joint approximate diagonalization (center), and in the basis obtained with FSPM (right); color coding is as in the top row. *Bottom*: Correspondence matrices in the three bases. Note how FSPM results in almost perfect alignment of basis functions.

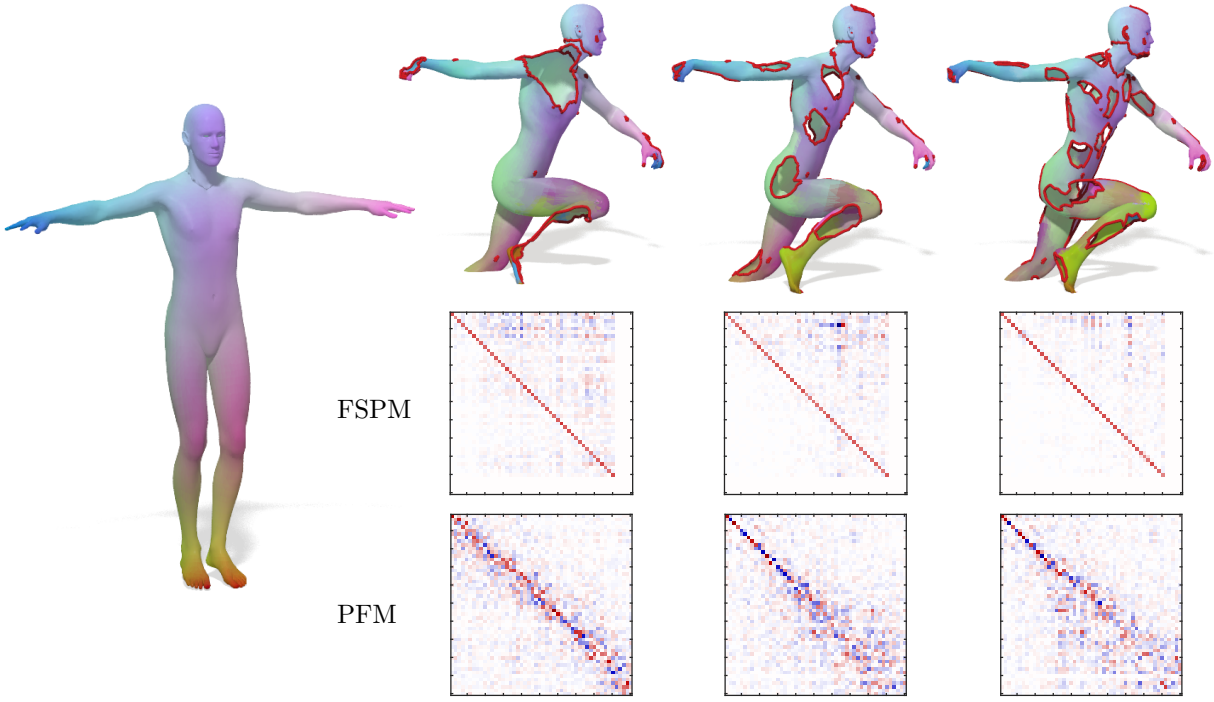

Figure 1.16: The three partial shapes shown above have *equal* missing area, although in different shapes and sizes. The resulting correspondence matrices in our new basis have same rank and similar off-diagonal patterns. Note how, differently from PFM, the diagonal in our representation remains sharp in all three cases.

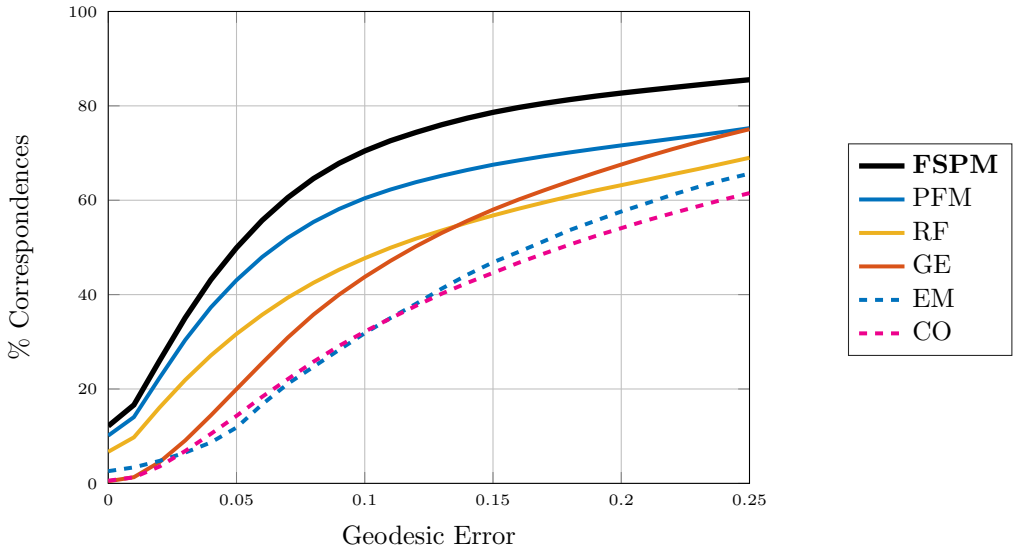

Figure 1.17: SHREC'16 Topology benchmark, includes shapes undergoing strong topological changes. While being fully spectral, FSPM improves upon PFM by over 10% at a fraction of the computational cost. Sparse correspondence methods (producing 250-1000 matches) and dense correspondence methods are denoted by dashed and solid lines, respectively.

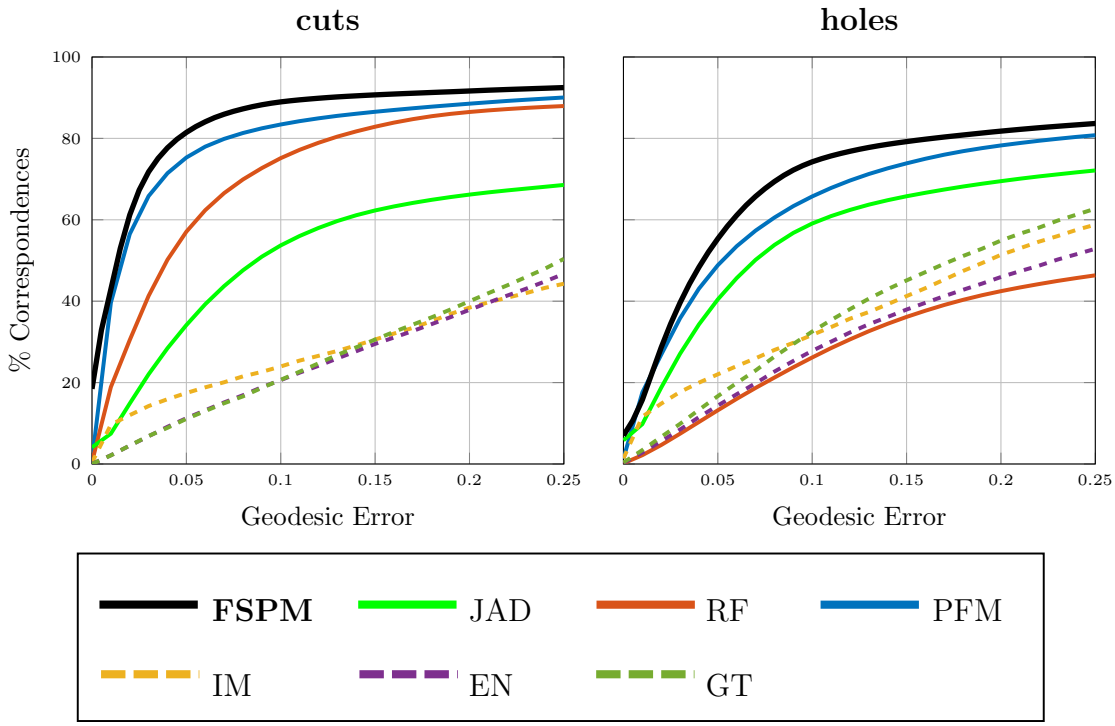

Figure 1.18: SHREC'16 Partial Correspondence benchmark. FSPM compares favorably with PFM, while being considerably more efficient (see also Figure 1.21). Sparse correspondence methods (producing 50-100 matches) and dense correspondence methods are denoted by dashed and solid lines, respectively.

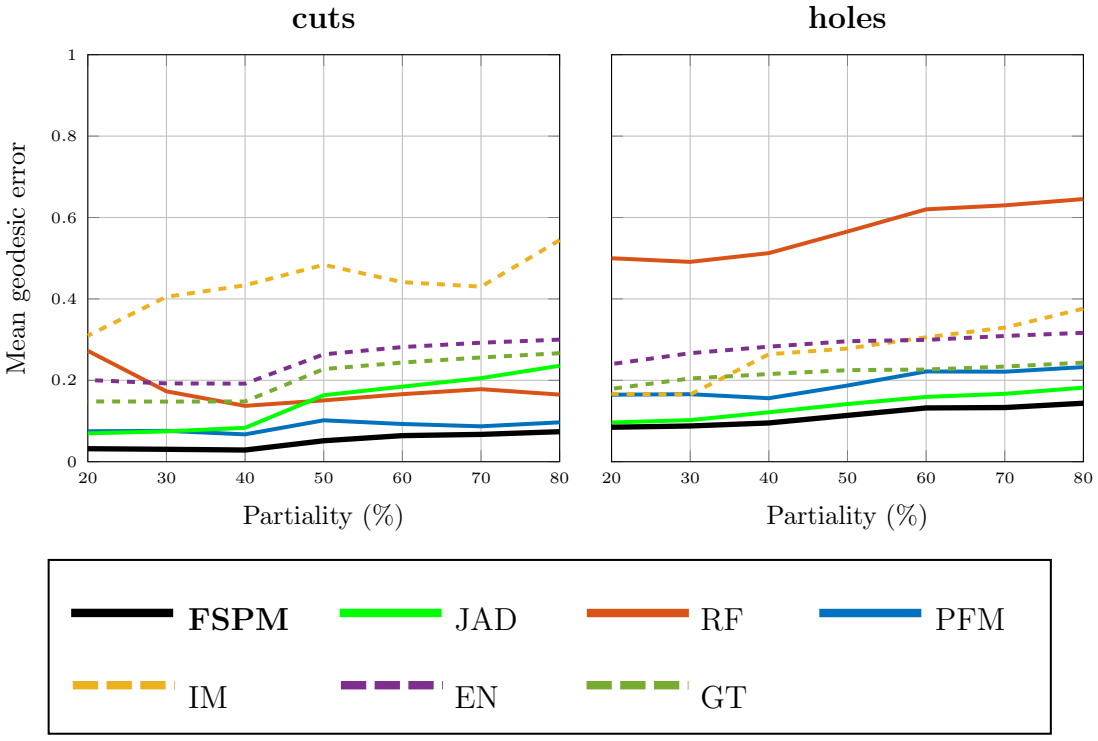

Figure 1.19: Performance of different methods on the SHREC'16 Partial Correspondence benchmark at increasing levels of partiality (measured as percentage of missing area).

Figure 1.20: Typical failure cases of FSPM. *Middle*: strong topological noise and partiality may affect the local descriptors, making the data term unreliable. *Right*: FSPM solutions are not guaranteed continuous, and may exhibit a mixture of inconsistently oriented patches due to symmetry ambiguity. For example, the hands of the kid are swapped and the belly is mapped to the back.

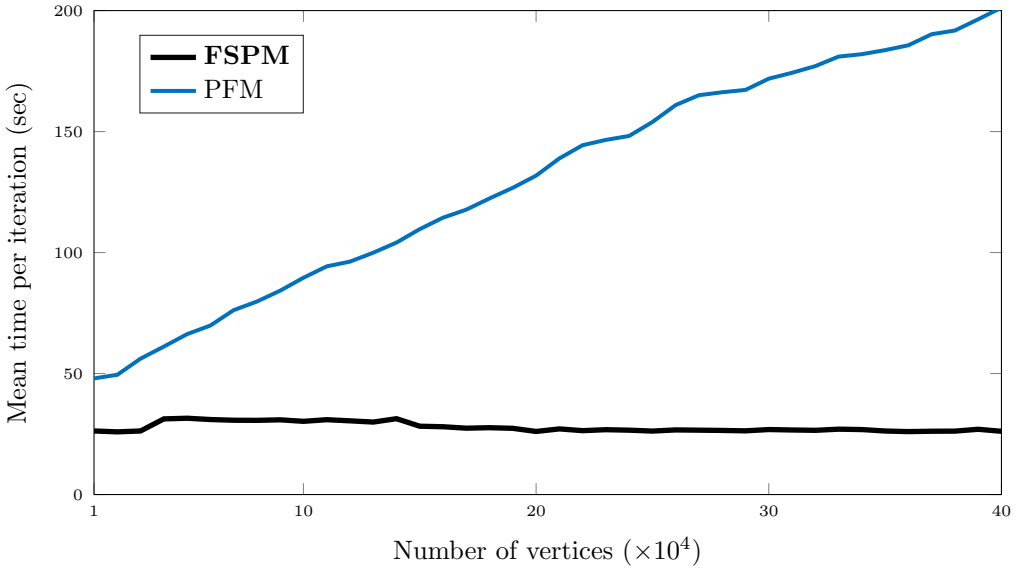

Figure 1.21: The runtime complexity of FSPM is constant w.r.t. shape size, while PFM exhibits a linear growth due to its explicit optimization in the spatial domain.

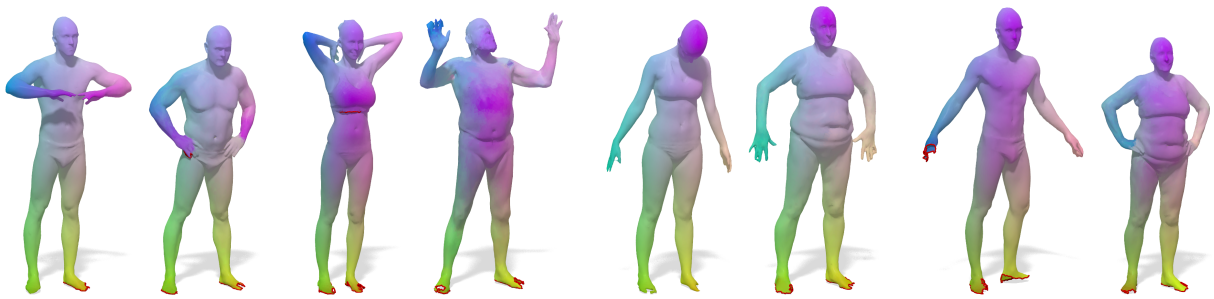

Figure 1.22: Example of solutions of FSPM on four pairs of shapes from the FAUST *real world* dataset [5]. The shapes contain several artifacts such as scanning noise, missing parts (on the feet) and topological merging due to self-contact (prominent in most examples).

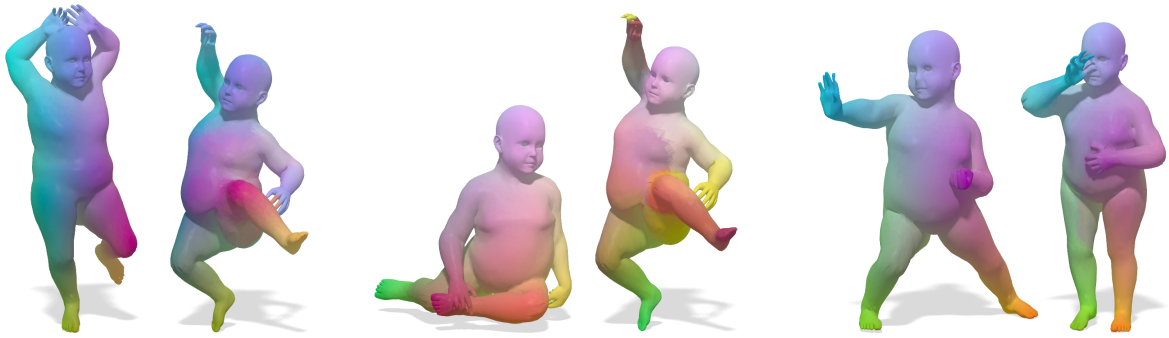

Figure 1.23: Examples of solutions obtained with FSPM in the presence of strong topological changes (SHREC'16 Topology benchmark [30]). Note how the quality of the correspondence remains largely unaffected even around the areas of contact.

Figure 1.24: Examples of dense partial shape correspondence obtained with FSPM on the SHREC'16 Partial Correspondence dataset [18]. The partial shapes are matched to the references shown on the left. Corresponding points have the same color; heat maps encode distance from the ground-truth (white denotes zero error, hot colors denote large error).

Bibliography

[1] Dror Aiger, Niloy J Mitra, and Daniel Cohen-Or. 4-points congruent sets for robust pairwise surface registration. *TOG*, 27(3):85, 2008.

[2] A. Albarelli, E. Rodolà, and A. Torsello. Fast and accurate surface alignment through an isometry-enforcing game. *Pattern Recognition*, 48:2209–2226, 2015.

[3] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1626–1633. IEEE, 2011.

[4] S Biasotti, A Cerri, A Bronstein, and M Bronstein. Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer Graphics Forum*, 2015.

[5] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR*, June 2014.

[6] Nicolas Boumal, Bamdev Mishra, P.-A. Absil, and Rodolphe Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.

[7] Alexander Bronstein, Michael Bronstein, Alfred Bruckstein, and Ron Kimmel. Partial similarity of objects, or how to compare a centaur to a horse. *IJCV*, 84(2):163–183, 2009.

[8] Alexander Bronstein, Michael Bronstein, and Ron Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.

[9] Alexander M Bronstein and Michael M Bronstein. Not only size matters: regularized partial matching of nonrigid shapes. In *Proc. NORDIA*, 2008.

[10] Alexander M Bronstein and Michael M Bronstein. Regularized partial matching of rigid shapes. In *Proc. ECCV*. 2008.

[11] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.

[12] Alan Brunton, Michael Wand, Stefanie Wuhrer, Hans-Peter Seidel, and Tino Weinkauf. A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models*, 76(2):70 – 85, 2014.

[13] Jean-Francois Cardoso and Antoine Souloumiac. Jacobi angles for simultaneous diagonalization. *SIAM J. Matrix Anal. Appl.*, 17(1):161–164, January 1996.

[14] Qifeng Chen and Vladlen Koltun. Robust nonrigid registration by convex optimization. In *Proc. ICCV*, 2015.

[15] L. Cosmo, E. Rodolà, A. Albarelli, F. Mémoli, and D. Cremers. Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum*, 2016.

[16] L. Cosmo, E. Rodolà, M. M. Bronstein, A. Torsello, D. Cremers, and Y. Sahillioğlu. Shrec'16: Partial matching of deformable shapes. In *Proc. 3DOR*, 2016.

[17] L. Cosmo, E. Rodolà, J. Masci, A. Torsello, and M. Bronstein. Matching deformable objects in clutter. In *Proc. 3DV*, 2016.

[18] Luca Cosmo, Emanuele Rodolà, Michael M. Bronstein, et al. SHREC'16: Partial matching of deformable shapes. In *Proc. 3DOR*, 2016.

[19] Mingsong Dou, Jonathan Taylor, Henry Fuchs, Andrew Fitzgibbon, and Shahram Izadi. 3D scanning deformable objects with a single RGBD sensor. In *Proc. CVPR*, 2015.

[20] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proc. SIGGRAPH*, pages 209–216, 1997.

[21] Qi-Xing Huang, Simon Flöry, Natasha Gelfand, Michael Hofer, and Helmut Pottmann. Reassembling fractured objects by geometric matching. *TOG*, 25(3):569–578, 2006.

[22] Qixing Huang, Fan Wang, and Leonidas J. Guibas. Functional map networks for analyzing and exploring large shape collections. *TOG*, 33(4):36, 2014.

[23] Varun Jain, Hao Zhang, and Oliver van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal of Shape Modeling*, 13(01):101–124, 2007.

[24] C. Kerl, J. Sturm, and D. Cremers. Dense visual slam for rgb-d cameras. In *Proc. IROS*, 2013.

[25] Vladimir G. Kim, Yaron Lipman, and Thomas A. Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011.

[26] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. *Computer Graphics Forum*, 32(2):439–448, 2013.

[27] A. Kovnatsky, K. Glashoff, and M. Bronstein. MADMM: a generic algorithm for non-smooth optimization on manifolds. In *Proc. ECCV*, 2016.

[28] Artiom Kovnatsky, Michael M. Bronstein, Xavier Bresson, and Pierre Vandergheynst. Functional correspondence by matrix completion. *CoRR*, abs/1412.8070, 2014.

[29] Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. Coupled quasi-harmonic bases. In *Computer Graphics Forum*, volume 32, pages 439–448, 2013.

[30] Zorah Lähner, Emanuele Rodolà, Michael M. Bronstein, et al. SHREC'16: Matching of deformable shapes with topological noise. In *Proc. 3DOR*, 2016.

[31] Hao Li, Robert W. Sumner, and Mark Pauly. Global correspondence optimization for non-rigid registration of depth scans. In *Proc. SGP*, pages 1421–1430, 2008.

[32] Or Litany, Alexander M Bronstein, and Michael M Bronstein. Putting the pieces together: Regularized multi-part shape matching. In *Proc. NORDIA*, 2012.

[33] Or Litany, Emanuele Rodolà, Alex Bronstein, and Michael Bronstein. Fully spectral partial shape matching. *Computer Graphics Forum*, 36(2):1681–1707, 2017.

[34] Or Litany, Emanuele Rodolà, Alex M. Bronstein, Michael M. Bronstein, and Daniel Cremers. Non-rigid puzzles. *Computer Graphics Forum*, 35(5):135–143, August 2016.

[35] J. Masci, D. Boscaini, M. M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proc. 3dRR*, 2015.

[36] Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008.

[37] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization&Mathematics*, pages 35–57, 2003.

[38] David Mumford and Jayant Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure and Applied Math.*, 42(5):577–685, 1989.

[39] Richard A Newcombe, Dieter Fox, and Steven M Seitz. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proc. CVPR*, 2015.

[40] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *Proc. ISMAR*, pages 127–136, 2011.

[41] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *Trans. Graphics*, 31(4):30, 2012.

[42] Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993.

[43] 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.

[44] E. Rodolà, A.M. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, pages 182–189, June 2012.

[45] E. Rodolà, S. Rota Bulò, and D. Cremers. Robust region detection via consensus segmentation of deformable shapes. *Computer Graphics Forum*, 33(5):97–106, 2014.

[46] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Computer Graphics Forum*, 2016.

[47] E. Rodolà, M. Moeller, and D. Cremers. Point-wise map recovery and refinement from functional correspondence. In *Proc. VMV*, 2015.

[48] 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, 2014.

[49] E. Rodolà, A. Torsello, T. Harada, Y. Kuniyoshi, and D. Cremers. Elastic net constraints for shape matching. In *Proc. ICCV*, pages 1169–1176, December 2013.

[50] Y Sahillioğlu and Y Yemez. Partial 3-d correspondence from shape extremities. *Computer Graphics Forum*, 33(6):63–76, 2014.

[51] Y. Sahillioğlu and Y. Yemez. Scale normalization for isometric shape matching. *Computer Graphics Forum (Proc. Pacific Graphics)*, 31(7), 2012.

[52] Yusuf Sahillioğlu and Yücel Yemez. Minimum-distortion isometric shape correspondence using EM algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(11):2203–2215, 2012.

[53] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP*, 2009.

[54] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *Proc. ECCV*, pages 356–369, 2010.

[55] Oliver van Kaick, Hao Zhang, and Ghassan Hamarneh. Bilateral maps for partial matching. *Comput. Graph. Forum*, 32(6):189–200, September 2013.

[56] Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.

[57] Luminita A Vese and Tony F Chan. A multiphase level set framework for image segmentation using the Mumford and Shah model. *IJCV*, 50(3):271–293, 2002.

[58] Linyu Wei, Qixing Huang, Duygu Ceylan, Etienne Vouga, and Hao Li. Dense human body correspondences using convolutional networks. *arXiv*, 1511.05904, 2015.

[59] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. Large-scale integer linear programming for orientation preserving 3d shape matching. *Computer Graphics Forum*, 30(5):1471–1480, 2011.