Fully Spectral Partial Shape Matching

O. Litany[1,4], E. Rodolà[2], A. M. Bronstein[1,3,4], M. M. Bronstein[1,2,4]

[1]Tel Aviv Univeristy, Israel   [2]University of Lugano, Switzerland   [3]Technion, Israel   [4]Intel Perceptual Computing, Israel
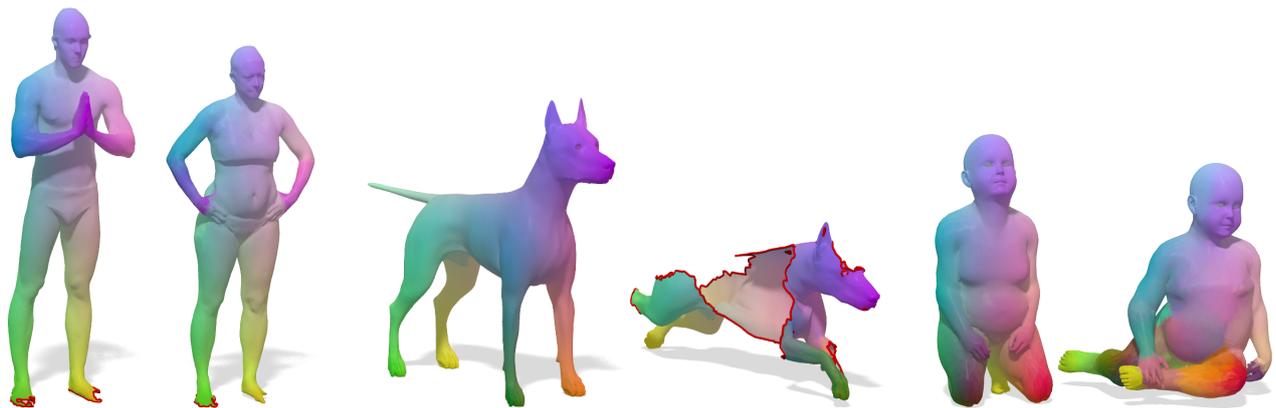
Figure 1: *Examples of dense correspondence computed with our method on real 3D scans (left pair, the areas of contact are glued together), missing parts (middle) and strong topological artifacts (right, touching parts are glued together). Corresponding points are encoded with the same color.*

Abstract

We propose an efficient procedure for calculating partial dense intrinsic correspondence between deformable shapes performed entirely in the spectral domain. Our technique relies on the recently introduced partial functional maps formalism and on the joint approximate diagonalization (JAD) of the Laplace-Beltrami operators previously introduced for matching non-isometric shapes. We show that a variant of the JAD problem with an appropriately modified coupling term (surprisingly) allows to construct quasi-harmonic bases localized on the latent corresponding parts. This circumvents the need to explicitly compute the unknown parts by means of the cumbersome alternating minimization used in the previous approaches, and allows performing all the calculations in the spectral domain with constant complexity independent of the number of shape vertices. We provide an extensive evaluation of the proposed technique on standard non-rigid correspondence benchmarks and show state-of-the-art performance in various settings, including partiality and the presence of topological noise.

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 3D shapes is one of the prototypical problems in computer graphics, geometric processing, and vision. Different flavors of this problem arise in applications ranging from texture and animation [vKZHCO11] to marker-less motion capture [WHC*16]. The recent progress in the development of commercial real-time 3D scanning technology has brought the need for fast, accurate, and reliable correspondence methods capable of dealing with real-world noise and artifacts. Particularly challenging settings of the correspondence problem include: non-rigid corre-

spondence, where the shapes are allowed to undergo deformations (in the simpler case assumed to be approximately isometric, and in the more difficult one, non-isometric); partial correspondence, where a subset of the shape has to be matched to its deformed full version; and geometric and topological noise (the latter arising, for example, due to occluded parts in the acquisition process). Practical scenarios involving real data acquisition consist of a combination of the above artifacts, where partiality, clutter, and topological noise are notoriously hard. The main goal of our paper is to deal with these challenging settings.

Related work. Classical non-rigid correspondence methods assume the shapes to be approximately isometric [BBK06] or topologically equivalent [LF09]. Such assumptions are usually broken in the partial correspondence and topological noise settings. Most recent works addressing the full correspondence problem [APL15, MDK*16, SPKS16, VLR*17] can cope with the lack of isometry to some extent, but cannot deal with significant partiality. *Rigid* partial correspondence has been addressed by regularized versions of the iterative closest point (ICP) algorithms [AMCO08, ART15]. Extensions of these methods to the non-rigid setting were proposed in [LSP08] but rely on an approximate initial alignment, practically limiting the application to small deformations arising, for example, in 3D tracking applications.

Bronstein *et al.* [BB08, BBBK09] proposed using a combined optimization problem over correspondence and parts, optimizing for metric distortion and the regularity of corresponding parts. In follow-up works, Rodolà *et al.* [RBA*12] relaxed the regularity requirement by allowing sparse correspondences and introduced a mechanism to explicitly control the degree of sparsity of the solution [RTH*13]. Sahillioğlu and Yemez [SY14] used a voting approach to match shape extremities, assuming them to be preserved by the partiality transformation. The common deficiencies of the above non-rigid partial correspondence methods are their ability to provide only a sparse correspondence (typically, order of tens of points), high computational complexity, and inability to deal with extreme partiality where boundary effects play a significant role.

Pokrass *et al.* [PBB13a] proposed a descriptor-based partial matching approach where the optimization over parts is done to maximize the matching of bags of local descriptors. The main drawback of this approach is that it only finds similar parts, without providing a correspondence between them. Windheuser *et al.* [WSSC11] formulated the shape matching problem as minimal surfaces on the product manifold, whose solution provides a guaranteed continuous and orientation-preserving matching. The method was shown to allow dealing with partiality; its main limitation is the assumption of mesh water-tightness and extreme computational complexity. Brunton *et al.* [BWW*14] computed partial correspondence through the alignment of tangent spaces, propagating some initial sparse correspondence. In the context of collections of shapes, partial correspondence has been considered in [VKTS*11, CGH14, CRA*16].

More closely related to our approach is the family of methods based on the notion of functional correspondence. Ovsjanikov *et al.* [OBCS*12] introduced *functional maps*, modeling correspondences as linear operators between spaces of functions on manifolds, which can be recovered provided a small set of known corresponding functions. In the Laplacian eigenbases, such operators can be efficiently approximated as low-rank matrices with approximately diagonal structure. Kovnatsky *et al.* [KBB*13] proposed finding bases that optimize for the diagonal structure of the correspondence matrix using joint diagonalization of Laplacians. Pokrass *et al.* [PBB*13b] extended functional maps to the setting where the ordering of the corresponding functions is unknown, solving simultaneously for a correspondence matrix and a permutation of the corresponding functions. Kovnatsky *et al.* [KBBV15] computed functional maps solving a geometric matrix completion problem, showing that non-isometric and mild partiality settings could be addressed to some extent. Huang *et al.* [HWG14] considered correspondence between collections of non-isometric shapes. Overall, while some formulations of functional maps allow to deal with the lack of isometry and partiality, this framework is in principle not designed to deal with partial correspondence.

Recently, Rodolà *et al.* [RCB*16] provided an empirical evidence and theoretical analysis of a surprising property of interaction between Laplacian eigenfunctions as the result of removing parts from surfaces. Applied to functional correspondence operators, this property results in a special slanted diagonal structure of the correspondence matrix. Based on this observation, the authors proposed the *partial functional maps* (PFM) framework. Follow-up works extended PFM to the cluttered non-rigid correspondence [CRM*16] and multiple shape ('non-rigid puzzles') [LRB*16b] settings. While showing impressive quality in challenging settings, the key deficiency of this line of works is an explicit model of parts, resulting in a complicated alternating optimization over the functional correspondence matrix represented in the spectral domain (w.r.t. the Laplacian eigenbasis) and the parts indicator function represented in the spatial domain (w.r.t. the Dirac basis).

Contribution. In this paper, we propose an efficient spectral-domain method for finding partial dense intrinsic correspondence between non-rigid shapes, which is capable of dealing with extreme partiality settings, topological noise, and non-isometric deformations. Our framework mainly builds on the previous works on partial functional maps [RCB*16] and joint approximate diagonalization (JAD) of Laplacians [KBB*13]. We show that a variant of the JAD problem with an appropriately modified coupling term allows to construct quasi-harmonic bases localized on the latent corresponding parts. This circumvents the need to explicitly compute the unknown parts by means of the cumbersome alternating minimization used in the partial functional maps approach, allowing to perform all the calculations *in the spectral domain* with constant complexity independent of the number of shape vertices.

Compared to the previous approaches our method offers a significant computational advantage, has an intuitive geometric interpretation, and is more robust to large amounts of partiality and lack of isometry. We show that the proposed approach achieves state-of-the-art performance on the most recent SHREC partial [CRB*16] and topological noise [LRB*16a] correspondence benchmarks.

2. Background

We model shapes as 2-manifolds $\mathcal{M}$ (possibly with boundary $\partial\mathcal{M}$) equipped with the area element $d\mu$ 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

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

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

with 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 $\phi_1, \phi_2, \dots$ 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 $L^2(\mathcal{M}) = \{f : \mathcal{M} \to \mathbb{R} \mid \int_{\mathcal{M}} f^2 d\mu < \infty\}$, 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

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

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

Functional correspondence. Our method builds upon the functional maps framework of Ovsjanikov *et al.* [OBCS*12]. 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

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

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 (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 [OBCS*12] has the general form of

$$\min_{\mathbf{C}} \|\mathbf{C}\mathbf{A} - \mathbf{B}\|_{\mathrm{F}}^2, \tag{5}$$

with the additional orthogonality constraint $\mathbf{C}^\top \mathbf{C} = \mathbf{I}$ if the underlying map is known to be area-preserving [OBCS*12].

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 [KBB*13] 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,

$$\hat{\phi}_i = \sum_{j=1}^k p_{ji} \phi_j, \qquad \hat{\psi}_i = \sum_{j=1}^k q_{ji} \psi_j \tag{6}$$

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 $\sum_{i=1}^k \langle \hat{\phi}_i, \Delta_{\mathcal{M}} \hat{\phi}_i \rangle_{\mathcal{M}} = \mathrm{tr}(\mathbf{P}^\top \mathbf{\Lambda}_{\mathcal{M}} \mathbf{P})$, 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 [CS96], arriving at the optimization problem

$$\min_{\mathbf{P}, \mathbf{Q}} \ \mathrm{off}(\mathbf{P}^\top \mathbf{\Lambda}_{\mathcal{M}} \mathbf{P}) + \mathrm{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{\mathrm{F}}^2 \tag{7}$$
$$\text{s.t.} \ \ \mathbf{P}^\top \mathbf{P} = \mathbf{I}, \quad \mathbf{Q}^\top \mathbf{Q} = \mathbf{I},$$

where $\mathrm{off}(\mathbf{A}) = \sum_{i \neq j} a_{ij}^2$. Problem (7) can be interpreted as a *joint approximate diagonalization* of the Laplacians $\Delta_{\mathcal{M}}$ and $\Delta_{\mathcal{N}}$ [KBB*13]. Note that if $\mu = 0$ (*i.e.*, no coupling) the global solution to (7) is $\mathbf{P} = \mathbf{Q} = \mathbf{I}$, resulting in the standard eigenfunctions of $\Delta_{\mathcal{M}}$ and $\Delta_{\mathcal{N}}$ when plugged into (6).

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 (7) by optimizing for a new basis on one shape only and keeping the other fixed to the standard Laplacian eigenfunctions,

$$\min_{\mathbf{Q} \in S(k,k)} \ \mathrm{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{\mathrm{F}}^2, \tag{8}$$

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 (7–8) are instances of *manifold optimization* and can be solved using efficient numerical techniques performing optimization on the matrix manifold [BMAS14].

Robust formulation. In practical settings, the corresponding functions f_i, g_i might be noisy, such that $Tf_i \neq g_i$ 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 (8) 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 achie-
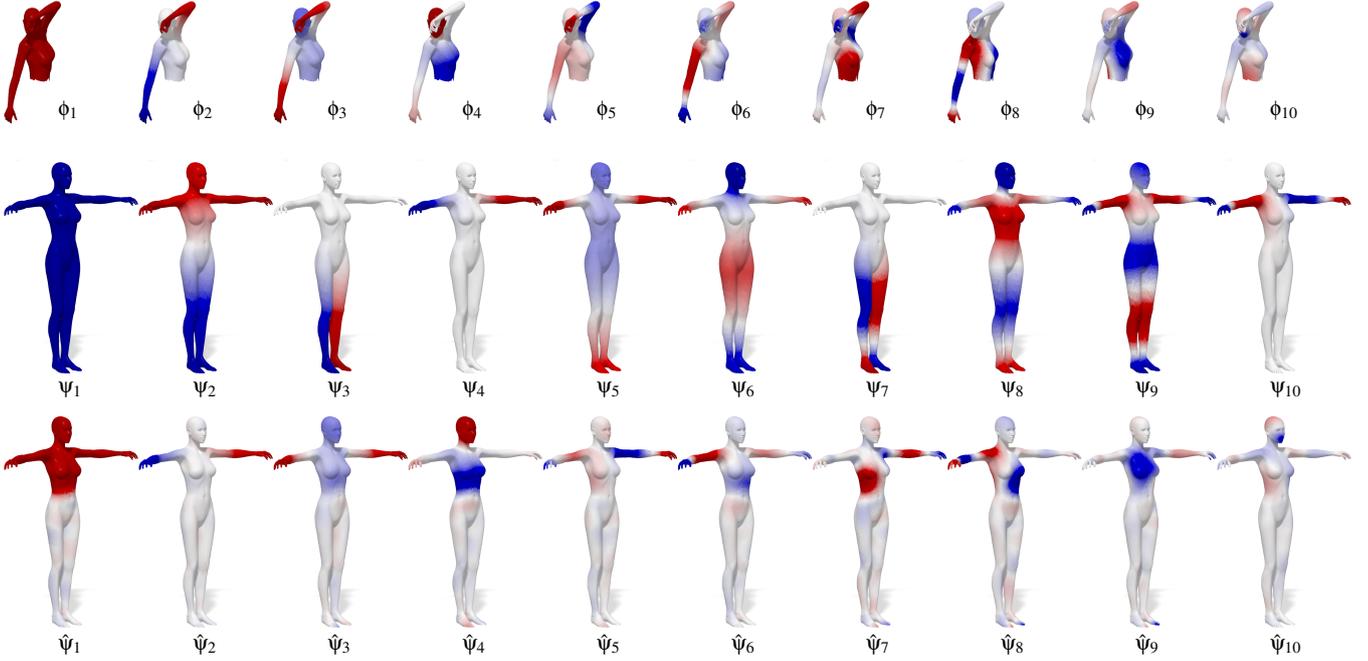
Figure 2: *The standard Laplacian eigenfunctions (two first rows) are strongly affected by the lack of perfect isometry and in the presence of missing parts. In the top and middle rows we show the first ten eigenfunctions $\{\phi_i\}_{i=1}^{10}$ and $\{\psi_j\}_{j=1}^{10}$ on a partial and full shape 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 our method: the new basis manifests the same behavior as in the first row, and is at the same time localized on the latent corresponding part.*

ved 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 [KGB16] can be employed to reach a good local optimum.

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 $\mathcal{N}'$.

Recently, Rodolà *et al.* [RCB*16] 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 2). 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 $\mathbf{\Delta}_{\mathcal{M}}$, $\mathbf{\Delta}_{\mathcal{N}}$ 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 $\mathbf{\Delta}_{\mathcal{N}}$ will assume the structure

$$\mathbf{\Delta}_{\mathcal{N}} = \begin{pmatrix} \mathbf{\Delta}_{\mathcal{M}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Delta}_{\bar{\mathcal{M}}} \end{pmatrix} + \begin{pmatrix} \mathbf{P}_{\mathcal{M}} & \mathbf{E} \\ \mathbf{E}^{\top} & \mathbf{P}_{\bar{\mathcal{M}}} \end{pmatrix}, \qquad (9)$$

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 $\partial\bar{\mathcal{M}}$.
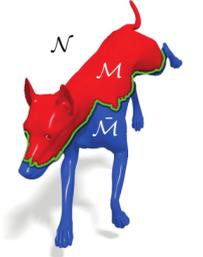
If the perturbation matrix is identically zero, then (9) is exactly block-diagonal; this describes the case in which $\mathcal{M}$ and $\bar{\mathcal{M}}$ are disjoint parts, and the eigenpairs of $\mathbf{\Delta}_{\mathcal{N}}$ are an interleaved sequence of those of the two blocks. The key result shown in [RCB*16] is that this interleaving property still holds even when considering the full matrix $\mathbf{\Delta}_{\mathcal{N}}$ as given in (9): Its eigenpairs consist of those of the blocks $\mathbf{\Delta}_{\mathcal{M}}$, $\mathbf{\Delta}_{\bar{\mathcal{M}}}$, up to some bounded 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 [RCB*16] has the form

$$\min_{\mathbf{C}, v} \|\mathbf{C}\mathbf{A} - \mathbf{B}(v)\| + \rho_{\mathrm{corr}}(\mathbf{C}) + \rho_{\mathrm{part}}(v), \qquad (10)$$

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 (10) 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 [RCB∗16] for the technical details). Note that function v is defined over the vertices of $\mathcal{N}$, hence scaling linearly with shape size. Problem (10) is optimized alternatingly over the Fourier and spatial domains in order to solve for correspondence and part respectively.

The above model was extended to multiple parts in [LRB∗16b] (a setting referred to as "non-rigid puzzle"), and in [CRM∗16] to deal with clutter. Despite the additional constraints on the parts, the problem formulations are essentially the same as in (10) and make use of similar regularizers on part and correspondence.

3. Fully spectral partial correspondence

The main drawback of partial functional maps [RCB∗16] and the follow-up works [LRB∗16b, CRM∗16] 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 5 for an evaluation). One of the main contributions of our paper is a simple observation allowing to formulate the partial functional maps problem *entirely in the spectral domain*. Our method bears resemblance to joint approximate diagonalization but has a fundamental difference that we will emphasize in the sequel.

Localization. A key feature of partial functional maps lies in their *spatially localized* behavior: Any solution to (10) 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 this paper we propose to absorb the spatial mask into a new basis $\{\hat{\psi}_j\}$ for $L^2(\mathcal{N})$; in doing so, we dispose of the explicit part v and obtain a simpler optimization problem, as we elucidate in the following.

Assume $\mathbf{C}, v$ are a solution to (10), 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

$$\sum_{ij}^{k} \langle f, \phi_i \rangle_{\mathcal{M}} c_{ji} \psi_j = \sum_{i=1}^{k} \langle v \cdot g, \psi_i \rangle_{\mathcal{N}} \psi_i \approx v \cdot g, \quad (11)$$

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

$$\sum_{ij}^{k} \langle f, \phi_i \rangle_{\mathcal{M}} \hat{\psi}_j \approx v \cdot g, \quad (12)$$

in other words, the modified basis $\{\hat{\psi}_j\}$ induces the sought localization. Importantly, in order for (12) 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 $\{\hat{\psi}_j\}$, we can phrase (12) in the spectral domain as:

$$\mathbf{A} \approx \mathbf{C}^{\top} \mathbf{B}(v) = \mathbf{C}^{\top} \mathbf{B}; \quad (13)$$

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

Problem. In light of our previous analysis, we consider the following manifold optimization problem:

$$\min_{\mathbf{Q} \in S(k,r)} \text{off}(\mathbf{Q}^{\top} \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{A}_r - \mathbf{Q}^{\top} \mathbf{B}\|_{2,1}, \quad (14)$$

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 (10) and can be estimated simply from the area ratio θ, or optimized for explicitly by solving (14) 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 [OBCS∗12, RCB∗16].

The optimization problem (14) 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 will allow us (see Eq. (16)) to tackle part-to-part settings as a simple modification to (14). The first r functions $\{\hat{\psi}_1, \ldots, \hat{\psi}_r\}$ 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 $\{\phi_1, \ldots, \phi_r\}$ under the functional correspondence (see Figure 2),

$$\langle T\phi_i, \hat{\psi}_j \rangle_{\mathcal{N}} \approx \delta_{ij}; \; i, j = 1, \ldots, r. \quad (15)$$

It is important to remark that, while the correct partial correspondence is a solution to our problem by Eq. (11–13), 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 5). In practice, 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 5). As observed by Litany *et al.* [LRB∗16b], 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

| Initialization | Iteration 5 | Iteration 25 | Iteration 75 | Iteration 150 | Iteration 700 | Iteration 1400 | Iteration 4000 |
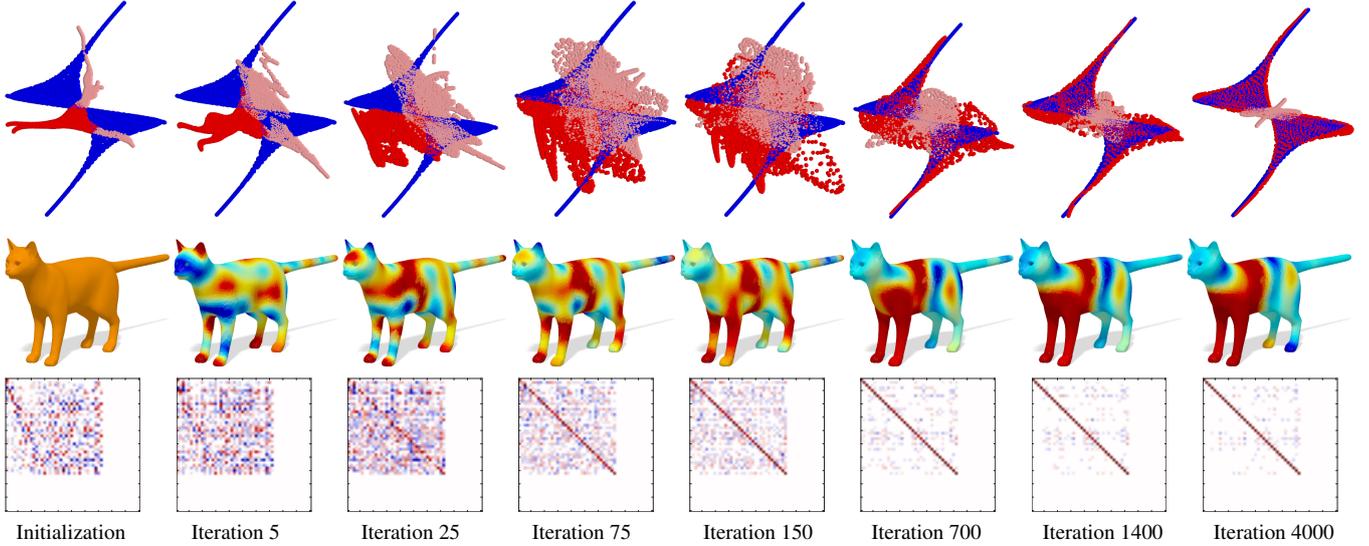
Figure 3: *Our optimization process admits an interpretation as a non-rigid alignment of the spectral embeddings of the input shapes (in this example, the cat meshes of Figure 7).* Top*: The spectral embeddings (in 2D for simplicity) at different time steps. The localization effect is manifested in the "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 our matching pipeline, and is being included here for illustration purposes.* Middle*: We plot the 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* $\mathbf{C}(t)$.

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 (10) 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 (14) we have to modify the coefficients $\mathbf{A}$ as well. This leads to

$$\min_{(\mathbf{P},\mathbf{Q})\in S^2(k,r)} \mathrm{off}(\mathbf{P}^\top\mathbf{\Lambda}_{\mathcal{M}}\mathbf{P}) + \mathrm{off}(\mathbf{Q}^\top\mathbf{\Lambda}_{\mathcal{N}}\mathbf{Q}) + \mu\|\mathbf{P}^\top\mathbf{A} - \mathbf{Q}^\top\mathbf{B}\|_{2,1},$$
(16)

where optimization is now performed on the *product of Stiefel manifolds*. In Figure 5 we illustrate 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 (14) and (16) can be viewed as variants of joint approximate diagonalization pro-

blems (8) and (7), respectively with the ℓ_2 data fitting term replaced by the more robust $\ell_{2,1}$ counterpart as was previously suggested in [PBB*13b] and [KGB16]. 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 (16) can be rewritten using full-rank
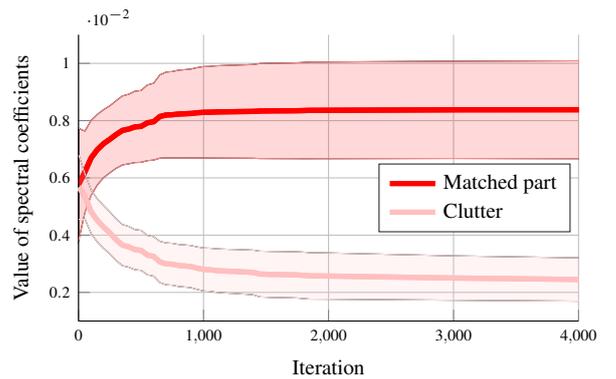
Figure 4: *This plot shows the evolution of the spectral coefficients of points belonging to the matchable (in red) and unmatchable (in pink) parts for the example shown in Figure 3. Observe the formation of two distinct groups of values from the very first iterations, with the values of the unmatched part tending towards zero.*
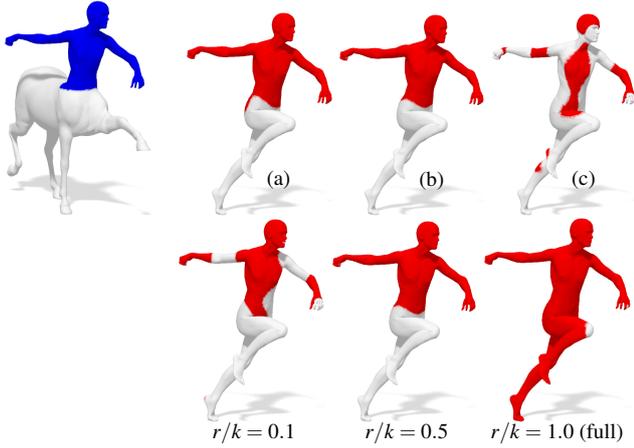
Figure 5: *Effect of changing the data term in* (16) *by using* (a) *descriptor fields localized to the correct region of the human shape (i.e., the region corresponding to the human part of the centaur);* (b) *descriptor fields supported on the entire shape, but similar only on the correct region; and* (c) *noisy descriptors with similar values outside the correct region. We visualize the segmentation of the human shape after optimization of* (16). *In the bottom row, we show localization as a function of the rank r.*

$k \times k$ 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 (7) with a modified metric. The effect of using the mask $\mathbf{W}_r$ is visualized in Figure 7. 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 (7) 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 (6) are aligned. Similarly, our approach (16) 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 our problem (16) 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 3, 4 and 7). Note that while some spectral approaches [JZvK07, MHK*08, RMC15] 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 (10) is also a solution to our formulation (14). A key difference lies in the direction of the map: If we regard matrix $\mathbf{Q}$ as the spectral representation of a functional map, our data term evaluates its pre-image in $\mathcal{M}$ (the partial shape), while (10) looks at the image on the full shape $\mathcal{N}$, thus requiring an explicit modeling of the part.

Further, as described in [RCB*16], the regularizer $\rho_{corr}(\mathbf{C})$ in

(10) includes a penalty term promoting $\mathbf{C}^\top \mathbf{C} \approx \begin{pmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. This area-preservation requirement is phrased as a hard constraint in our problem, where we optimize over the Stiefel manifold $S(k,r)$, such that $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_r$. Overall, the optimization problem (14) is less engineered than (10), has less parameters, and is simpler to optimize. See Figures 6-8 for further comparisons.

4. Implementation

Our optimization problem (14) is manifold-constrained and non-smooth (due to the $\ell_{2,1}$ norm). We solve it by using the MADMM scheme of Kovnatsky *et al.* [KGB16]. We implemented our method in Matlab using `manopt` [BMAS14], a framework for optimization over manifolds. Laplacians were discretized using the classical cotangent scheme [PP93, MDSB03]; 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. We further note that, although the manifold constraints render the problem non-convex and MADMM gives no global optimality guarantees, in practice we observed a stable behavior with a strictly decreasing cost value and fast convergence (an empirical evaluation is provided in Section 5).

Initialization. We initialize the orthogonal matrices in problems (14) and (16) 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 we avoid using this input by replacing $\{f_i, g_i\}_{i=1}^q$ with dense descriptor fields calculated on $\mathcal{M}$ and $\mathcal{N}$, where q is the number of dimensions of the descriptor. In all our experiments we used the 352-dimensional SHOT [TSDS10] with default parameters.

Point-wise map conversion and refinement. After convergence,
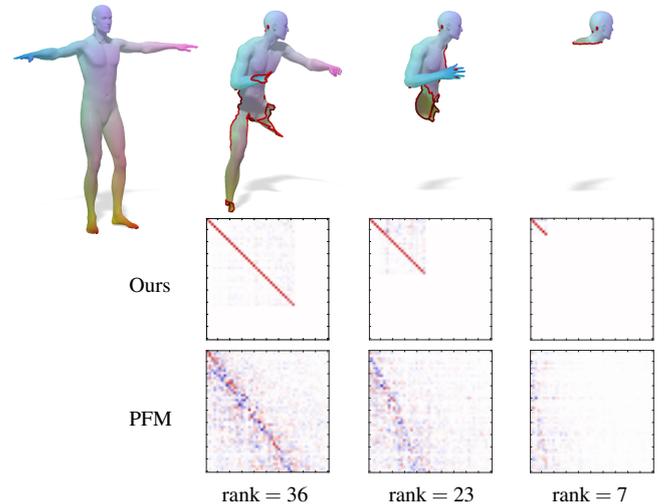
Figure 6: *Correspondence matrices at increasing partiality. We show the solutions obtained by our method in the new basis (middle row) and by PFM in the standard Laplacian eigenbasis (bottom row). Observe how our representation remains crisp even at extreme levels of partiality (rightmost column).*
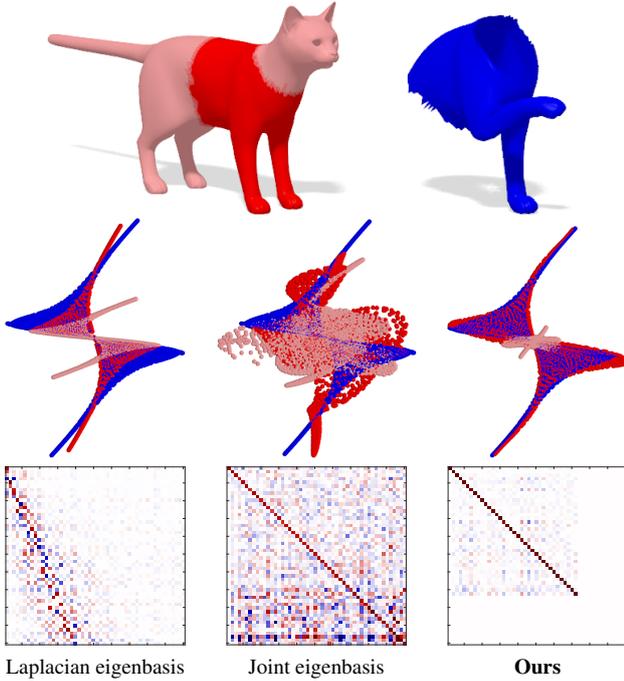
Figure 7: Top*: Input shapes (dark red denotes the part corresponding to the partial blue shape).* Middle*: Spectral embeddings (shown are the first two eigenfunctions) using the standard Laplacian eigenbasis (left), in the basis obtained by joint approximate diagonalization (center), and in the basis obtained with our approach (right); color coding is as in the top row.* Bottom*: Correspondence matrices in the three bases. Note how our method results in almost perfect alignment of basis functions.*

Figure 8: *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 9: *SHREC'16 Topology benchmark, includes shapes undergoing strong topological changes. While being fully spectral, our method 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.*

we recover the point-wise correspondence by a simple nearest-neighbor search in the k-dimensional spectral domain [OBCS*12] (*e.g.*, in the example of Figure 7, 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. We repeat this process 5 times. For a fair comparison, we applied the same refinement procedure to partial functional maps [RCB*16] and joint diagonalization [KBB*13]. Note that while more sophisticated recovery methods exist [RMC15], these work under the assumption of no partiality, which is violated in our setting.

5. Experimental results

We evaluated our method extensively in a variety of settings. Our method 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 [KLF11]. 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

$$\varepsilon(x) = \frac{d_{\mathcal{N}}(y,y^*)}{|\mathcal{N}|^{1/2}}, \qquad (17)$$

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

Topological changes We performed a full quantitative evaluation on the recent SHREC'16 Topology benchmark [LRB*16a] (*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 15). The methods appearing in the original benchmark are random forests (RF) [RRBW*14], Green's embedding (GE) [LRB*16a], and isometric embedding (EM) [SY12a]. As reported in Figure 9, previous approaches demonstrated poor performance
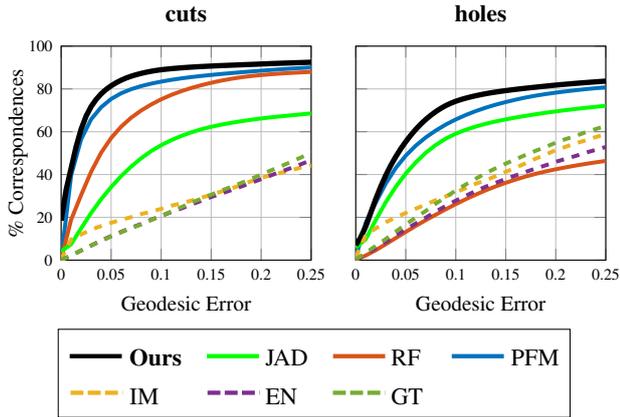
Figure 10: *SHREC'16 Partial Correspondence benchmark. Our method compares favorably with PFM, while being considerably more efficient (see also Figure 13). Sparse correspondence methods (producing 50-100 matches) and dense correspondence methods are denoted by dashed and solid lines, respectively.*

Figure 11: *Performance of different methods on the SHREC'16 Partial Correspondence benchmark at increasing levels of partiality (measured as percentage of missing area).*

due to the challenging setting. We additionally included in the comparison partial functional maps (PFM) [RCB*16] and the recent convex optimization (CO) method of Chen and Koltun [CK15], which performs an explicit modeling of topological artifacts but did not previously appear in the benchmark. For a fair comparison, we disabled the extrinsic regularization term of [CK15] since it relies on the shapes being approximately aligned in $\mathbb{R}^3$.

Part-to-full. We quantitatively evaluated our method in the partial matching scenario on the challenging SHREC'16 Partial Correspondence benchmark [CRB*16]. 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 10 and 11, and qualitatively in Figure 16. We compare with partial functional maps (PFM) [RCB*16], random forests (RF) [RRBW*14], scale-invariant isometric matching (IM) [SY12b], game-theoretic matching (GT) [RBA*12], and elastic net matching (EN) [RTH*13], as these methods appeared in the original benchmark. We additionally include joint diagonalization (JAD) [KBB*13] in the comparison.

We see from the plots that our method has a ∼10% improvement on PFM, the closest competitor, in both datasets. Given the purely spectral nature of our method 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 5) The poor performance of JAD puts in evidence the importance of correctly estimating the rank of the new basis, as we discussed in Section 3.
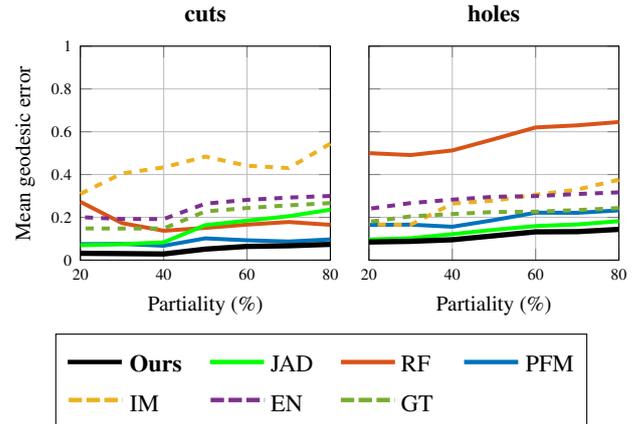
Scanned data. We carried out qualitative experiments on the FAUST dataset [BRLB14], which contains real human shapes acquired with a 3D scanning device. By nature of the acquisition process, these shapes are affected by topological artifacts as well as missing parts due to self-occlusions, resulting in a challenging testbed for shape matching. The results are shown in Figure 14.

Runtime. In Figure 13 we report a runtime comparison with PFM at increasing number of vertices. Since our 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 (14)). In contrast, due to the alternating optimization over the spectral and spatial domains, the runtime complexity of PFM grows linearly with shape size (Equation (10)). The average runtime on the SHREC'16 benchmarks was ∼220 sec. for our method and ∼1240 sec. for PFM.

6. Discussion and conclusions

We introduced a novel method for partial dense intrinsic correspondence between deformable shapes. Contrarily to existing approaches, our method 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 our method 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. Our method 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 12. The main limitation of our method 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 13: *The runtime complexity of our method is constant w.r.t. shape size, while PFM exhibits a linear growth due to its explicit optimization in the spatial domain.*

References

[AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust pairwise surface registration. *Trans. Graphics 27*, 3 (2008), 85. 2

[APL15] AIGERMAN N., PORANNE R., LIPMAN Y.: Seamless surface mappings. *ACM Trans. Graph. 34*, 4 (July 2015), 72:1–72:13. 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

[BB08] BRONSTEIN A. M., BRONSTEIN M. M.: Not only size matters: regularized partial matching of nonrigid shapes. In *Proc. NORDIA* (2008). 2

[BBBK09] 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

[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. 3, 7

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (June 2014). 9, 11

[BWW*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

[CGH14] CHEN Y., GUIBAS L., HUANG Q.: Near-optimal joint object matching via convex relaxation. In *Proc. ICML* (2014), pp. 100–108. 2

[CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *Proc. ICCV* (2015). 9

[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., ET AL.: SHREC'16: Partial matching of deformable shapes. In *Proc. 3DOR* (2016). 2, 9, 12

[CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSELLO A., BRONSTEIN M.: Matching deformable objects in clutter. In *Proc. 3DV* (2016). 2, 5

[CS96] CARDOSO J.-F., SOULOUMIAC A.: Jacobi angles for simultaneous diagonalization. *SIAM J. Matrix Anal. Appl. 17*, 1 (Jan. 1996), 161–164. 3

[HWG14] HUANG Q., WANG F., GUIBAS L. J.: Functional map networks for analyzing and exploring large shape collections. *TOG 33*, 4 (2014), 36. 2

[JZvK07] JAIN V., ZHANG H., VAN KAICK O.: Non-rigid spectral correspondence of triangle meshes. *International Journal of Shape Modeling 13*, 01 (2007), 101–124. 7

[KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. In *Computer Graphics Forum* (2013), vol. 32, pp. 439–448. 2, 3, 8, 9

[KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proc. CVPR* (2015). 2

[KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M.: MADMM: a generic algorithm for non-smooth optimization on manifolds. In *Proc. ECCV* (2016). 4, 6, 7

Figure 12: *Typical failure cases of our method.* Middle*: strong topological noise and partiality may affect the local descriptors, making the data term unreliable.* Right*: our 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 14: *Example of solutions of our method on four pairs of shapes from the FAUST real world dataset [BRLB14]. 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).*
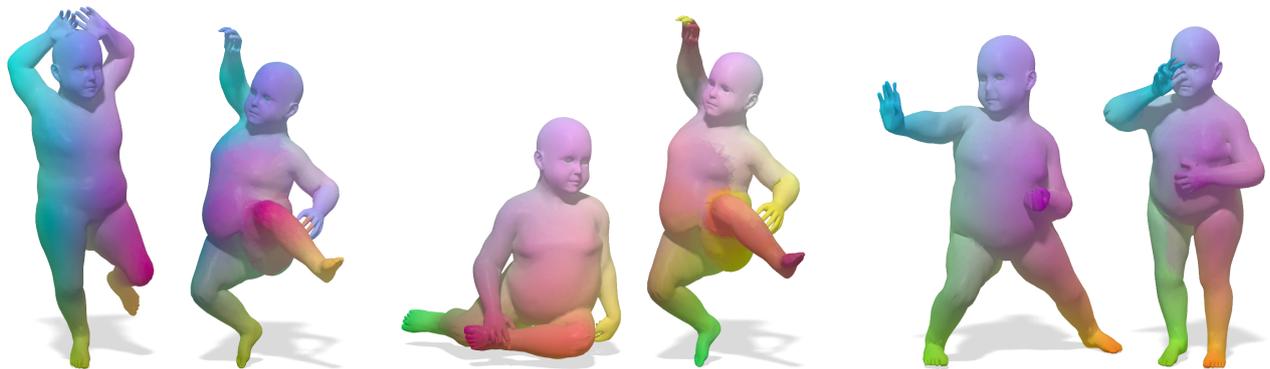
Figure 15: *Examples of solutions obtained with our method in the presence of strong topological changes (SHREC'16 Topology benchmark [LRB*16a]). Note how the quality of the correspondence remains largely unaffected even around the areas of contact.*

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG 30*, 4 (2011), 79. 8

[LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *Trans. Graphics 28*, 3 (2009). 2

[LRB*16a] LÄHNER Z., RODOLÀ E., BRONSTEIN M. M., ET AL.: SHREC'16: Matching of deformable shapes with topological noise. In *Proc. 3DOR* (2016). 2, 8, 11

[LRB*16b] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. *Computer Graphics Forum 35*, 5 (August 2016), 135–143. 2, 5

[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

[MDK*16] MARON H., DYM N., KEZURER I., KOVALSKY S., LIPMAN Y.: Point registration via efficient convex relaxation. *ACM Trans. Graph. 35*, 4 (July 2016), 73:1–73:12. 2

[MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization&Mathematics* (2003), 35–57. 7

[MHK*08] MATEUS D., HORAUD R. P., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR* (2008). 7

[OBCS*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, 5, 8

[PBB13a] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M.: Partial shape matching without point-wise correspondence. *Numer. Math. Theor. Meth. Appl. 6* (2013), 223–244. 2

[PBB*13b] 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. 2, 6

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental math. 2*, 1 (1993), 15–36. 7

[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, 9

[RCB*16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* (2016). 2, 4, 5, 7, 8, 9

[RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Proc. VMV* (2015). 7, 8

[RRBW*14] RODOLÀ E., ROTA BULÒ S., WINDHEUSER T., VESTNER M., CREMERS D.: Dense non-rigid shape correspondence using random forests. In *Proc. CVPR* (2014), pp. 4177–4184. 8, 9

[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, 9

[SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Trans. Graph. 35*, 4 (July 2016), 72:1–72:13. 2

[SY12a] SAHILLIOĞLU Y., YEMEZ Y.: Minimum-distortion isometric

Figure 16: *Examples of dense partial shape correspondence obtained with our method on the SHREC'16 Partial Correspondence dataset [CRB*16]. 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).*

shape correspondence using EM algorithm. *IEEE Trans. Pattern Anal. Mach. Intell. 34*, 11 (2012), 2203–2215. 8

[SY12b] SAHILLIOĞLU Y., YEMEZ Y.: Scale normalization for isometric shape matching. *Computer Graphics Forum (Proc. Pacific Graphics) 31*, 7 (2012). 9

[SY14] SAHILLIOĞLU Y., YEMEZ Y.: Multiple shape correspondence by dynamic programming. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 121–130. 2

[TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proc. ECCV* (2010), pp. 356–369. 7

[VKTS*11] VAN KAICK O., TAGLIASACCHI A., SIDI O., ZHANG H., COHEN-OR D., WOLF L., HAMARNEH G.: Prior knowledge for part correspondence. *Computer Graphics Forum 30*, 2 (2011), 553–562. 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

[VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. *arXiv:1701.00669* (2017). 2

[WHC*16] WEI L., HUANG Q., CEYLAN D., VOUGA E., LI H.: Dense human body correspondences using convolutional networks. In *Proc. CVPR* (2016). 1

[WSSC11] WINDHEUSER T., SCHLICKEWEI 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