Regularized Point-wise Map Recovery
from Functional Correspondence

E. Rodolà[1,3], M. Moeller[2,3], D. Cremers[3]

[1]USI Lugano, Switzerland   [2]University of Siegen, Germany   [3]TU Munich, Germany

Abstract

The concept of using functional maps for representing dense correspondences between deformable shapes has proven to be extremely effective in many applications. However, despite the impact of this framework, the problem of recovering the point-to-point correspondence from a given functional map has received surprisingly little interest. In this paper we analyze the aforementioned problem and propose a novel method for reconstructing point-wise correspondences from a given functional map. The proposed algorithm phrases the matching problem as a regularized alignment problem of the spectral embeddings of the two shapes. Opposed to established methods our approach does not require the input shapes to be nearly-isometric, and easily extends to recovering the point-to-point correspondence in part-to-whole shape matching problems. Our numerical experiments demonstrate that the proposed approach leads to a significant improvement in accuracy in several challenging cases.

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

1. Introduction

While identifying the similarity and correspondence of two shapes can be done quite accurately by a human observer, the automation of this shape matching task remains a challenging problem in many computer vision and graphics applications, see [KZHCO11] for an overview. The introduction of *functional correspondences* by Ovsjanikov *et al.* in [OBCS*12] is one of the most influential recent advances in the case where the two shapes are related by a non-rigid deformation. Instead of finding point-to-point mappings between the two shapes the authors proposed to determine a linear operator, the so-called *functional map*, that maps between the spaces of square integrable functions on the respective shapes. The recovery of a point-wise correspondence can be seen as the limiting case of mapping peaks on one shape to peaks on the other. The key advantage of the functional map idea is that it can be represented quite compactly as a low-dimensional matrix under a proper choice of basis functions. In [OBCS*12] the authors suggest the use of the Laplacian eigenfunctions as a natural basis for smooth functions on the respective shapes. Additionally, Laplacian eigenfunctions have the desirable property of being invariant to isometric transformations. Since smooth functions are well approximated by the first few Laplacian eigenfunctions,


Figure 1: *Given a functional map as input, our method allows to accurately recover and refine the underlying point-to-point mapping, even under non-isometric deformations. In the first row, color encodes distance to the ground-truth, increasing from white to red. The input map and its optimized version are shown in the second row.*

one can "truncate" the representation of the functional map by setting the images of all but the first few basis functions to

zero. This leads to a very low dimensional representation of the functional map while still keeping a good approximation to the true shape correspondence.

Extensions of the functional map framework have been proposed by several authors, covering the problem of non-isometric deformations [PBB*13, KBB*13, RBW*14, KBBV15], partial similarity [RCB*16, LRB*16], clutter [CRM*16], shape exploration [ROA*13, HWG14] and image segmentation [WHG13] among others. However, despite the success of these methods, there has been a general lack of interest on the inverse problem of accurately *reconstructing* a point-wise map from its functional representation – a common requirement in many practical applications.

The approach established by Ovsjanikov *et al.* [OBCS*12] recovers a point-wise map by a nearest-neighbor search in the embedded functional space. While this technique works well if the given functional map is sufficiently close to a point-wise map, its performance decreases significantly as the number of basis functions for the functional map is reduced. An iterative technique that alternates between estimating point-to-point correspondences and updating the functional map can improve the matching accuracy, however, only works under specific assumptions on the initial functional map and the type of deformation between the shapes. Extensions of such iterative techniques have recently been applied to retrieve exact bijections [VLR*17], for near-isometric partial matching [RCB*16], and in a correspondence-less setting for shape retrieval tasks [GT15].

Taking a slightly different direction of research, [KBBV15] proposed to evaluate the quality of a functional map *without* recovering a point-to-point correspondence by adopting a soft error criterion.

Contribution. This paper builds upon our previous work [RMC15] on the problem of accurate point-wise map recovery from a given functional map. The following aspects are our key contributions:

- To the best knowledge of the authors, we are the first to rigorously state and analyze the point-wise map recovery problem.
- We propose to consider the point-wise map recovery problem as a point cloud alignment problem in the embedded functional space, and use a regularized probabilistic model to ensure locally consistent point matching. Most importantly, our model is not based on the (common) assumption of the two shapes being related by a near-isometry, and therefore extends to a significantly larger class of possible applications.

We extend the results from our conference publication [RMC15] in several different ways, including:

- The extension of the proposed technique to the part-to-full shape matching problem.

- Extended numerical experiments further reinforcing our claim that our method significantly outperforms the existing methods in both, the nearly-isometric and the inter-class settings.
- Improving the applicability to large shape matching problems by providing an efficient gpu implementation of the proposed approach.

2. Background

We model shapes as two-dimensional Riemannian manifolds $\mathcal{M}$ (possibly with boundary) with area element $d\mu$. We denote the space of square-integrable functions on the manifold $\mathcal{M}$ by $L^2(\mathcal{M}) = \{f : \mathcal{M} \to \mathbb{R} \mid \int_{\mathcal{M}} f^2 d\mu < \infty\}$, and use the standard $L^2(\mathcal{M})$ inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} fg d\mu$. In analogy to the Laplace operator in flat spaces, the symmetric Laplace-Beltrami operator $\Delta_{\mathcal{M}}$ provides us with all necessary tools for extending Fourier analysis to manifolds. In particular, the manifold Laplacian yields an eigen-decomposition $\Delta_{\mathcal{M}} \phi_i = \lambda_i \phi_i$ for $i \geq 1$, with eigenvalues $0 = \lambda_1 < \lambda_2 \leq \dots$ and eigenfunctions $\{\phi_i\}_{i \geq 1}$ forming an orthonormal basis of $L^2(\mathcal{M})$. Due to this property, any function $f \in L^2(\mathcal{M})$ can be represented via the (manifold) Fourier series expansion

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

Functional correspondence. Consider two manifolds $\mathcal{M}$ and $\mathcal{N}$, and let $T : \mathcal{M} \to \mathcal{N}$ be a bijective mapping between them. While classical shape matching approaches try to identify point-to-point correspondences, *i.e.*, the bijection T, directly, the idea of functional maps proposed by Ovsjanikov *et al.* [OBCS*12] is to consider an operator $T_F : L^2(\mathcal{M}) \to L^2(\mathcal{N})$, mapping functions on $\mathcal{M}$ to functions on $\mathcal{N}$ via the composition $T_F(f) = f \circ T^{-1}$. It is remarkable that this seemingly simple idea allows to move from identifying a map between manifolds to identifying a linear operator between Hilbert spaces.

Because T_F is a linear operator, it is fully defined by the images of any basis $\{\phi_i\}_{i \geq 1}$ of $L^2(\mathcal{M})$. By also choosing a basis $\{\psi_i\}_{i \geq 1}$ of $L^2(\mathcal{N})$, the operator T_F can equivalently be represented by coefficients (c_{ij}), *e.g.*, for orthonormal $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ and arbitrary $f \in L^2(\mathcal{M})$ one finds

$$T_F(f) = \sum_{ij \geq 1} \langle f, \phi_j \rangle_{\mathcal{M}} \underbrace{\langle T_F(\phi_j), \psi_i \rangle_{\mathcal{N}}}_{c_{ij}} \psi_i. \qquad (2)$$

The application of T_F is expressed by linearly transforming the expansion coefficients of f from basis $\{\phi_i\}_{i \geq 1}$ onto basis $\{\psi_i\}_{i \geq 1}$. Since the transformation is encoded in the coefficients c_{ij}, the (possibly infinite) matrix $\mathbf{C} = (c_{ij})$ provides a representation of T_F for given bases. Seeking a functional correspondence among the two shapes then amounts to determining the unknown $\mathbf{C}$ that satisfies certain mapping constraints (*e.g.*, preservation of descriptors) [OBCS*12]
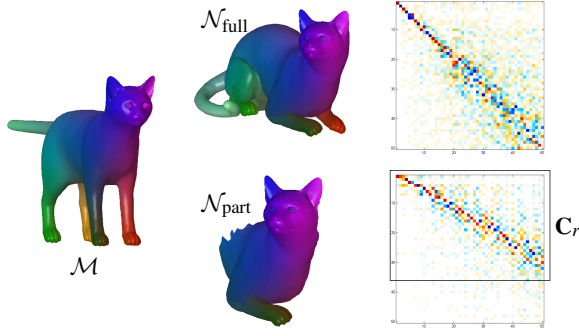
Figure 2: *The matrix representation of a functional map w.r.t. harmonic basis functions has a diagonal structure, with an angle depending on the amount of surface overlap between the two shapes. Note how the map between $\mathcal{M}$ and $\mathcal{N}_{part}$ (bottom) is manifestly low rank, with a rectangular sub-matrix $\mathbf{C}_r$ and a bottom block consisting of zeros. On the left, corresponding points are assigned the same color.*

Figure 3: *Each eigenfunction ψ_i of the Laplace-Beltrami operator of a partial shape (top row) can be put into approximate correspondence with an eigenfunction ϕ_j of a near-isometric full shape (bottom row) for some $j \geq i$. The vice versa does not hold.*

or manifesting a specific structure (*e.g.*, diagonally dominant) [PBB*13, KBBV15].

Remark. When choosing indicator vectors supported at each point as the finite bases $\{\phi_i\}_{i=1}^n$, $\{\psi_j\}_{j=1}^m$, one obtains the classical representation of the correspondence as a $n \times m$ binary matrix (a permutation matrix if $n = m$ and T is assumed to be bijective).

Basis truncation. As suggested in [OBCS*12], the eigenfunctions $\{\phi_i\}_{i\geq1}$, $\{\psi_i\}_{i\geq1}$ of the respective Laplacians yield a particularly interesting choice for the aforementioned bases: In case the two shapes to be compared are related by a near-isometry, the equality $\psi_i = \pm\phi_i \circ T^{-1}$ holds (approximately) for all $i \geq 1$, which leads to the matrix representation $\mathbf{C}$ of the functional map being diagonal, $c_{ij} = 0$ if $i \neq j$.

With this choice of a basis, the authors of [OBCS*12] proposed to truncate the series (2) 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. The matrix is diagonally dominant with an off-diagonal spread at the high end of the spectrum (see Fig. 2 top). Such a structure is manifested due to the approximate equality $c_{ij} = \langle T_F(\phi_j), \psi_i \rangle_{\mathcal{N}} \approx \pm\delta_{ij}$.

Partial functional correspondence. In case one of the two shapes has holes or missing parts, the functional representation of the correspondence still has a meaningful structure, as recently shown in [RCB*16]. Let $\mathcal{N}$ be a partial shape, $\mathcal{M}$ a full shape, and let $\mathcal{M}' \subset \mathcal{M}$ be the region of $\mathcal{M}$ corresponding to $\mathcal{N}$ under a near-isometry $T : \mathcal{N} \to \mathcal{M}'$. Then, for each eigenfunction ψ_i of $\mathcal{N}$ there exists an eigenfunction ϕ_j of $\mathcal{M}$ for some $j \geq i$, such that $\phi_j \approx \pm\psi_i \circ T^{-1}$ (see Fig. 3). In other words, the eigenfunctions of the Laplacian are still compatible under partiality, but some eigenfunctions
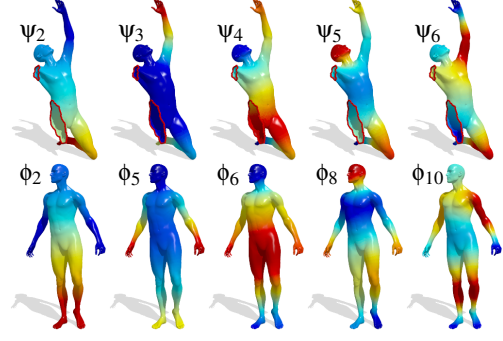
of the full shape do not have a corresponding counterpart on the partial shape. This results in a matrix $\mathbf{C}$ manifesting a slanted diagonal structure (Fig. 2 bottom), with an angle depending on the area ratio of the two surfaces [RCB*16]. Note that in this paper we only consider the part-to-whole case as done in [RCB*16, CRB*16], *i.e.*, only one of the two shapes is allowed to have missing parts.

3. Point-wise map recovery

The functional map representation greatly simplifies correspondence-based tasks. If the harmonic basis is truncated to k basis functions, the shape matching problem boils down to solving for k^2 unknowns, where k is possibly very small (typically in the range of few tens / hundreds) . At the same time, the truncation has the effect of 'low-pass' filtering, thus producing smooth correspondences. In many applications, however, it is desirable to reconstruct the point-to-point mapping induced by the functional map. Thus, the interest shifts to the *inverse* problem of recovering the map T from its functional representation T_F.

Let us now consider the discretized problem. Let $\mathcal{M}$ and $\mathcal{N}$ be represented by triangular meshes with m and n nodes respectively (with $\mathcal{M}$ being possibly a partial shape), and assume that the matrices $\Phi \in \mathbb{R}^{m \times k}$, $\Psi \in \mathbb{R}^{n \times k}$ contain the first k basis vectors of some ordered basis, *e.g.*, arising from the eigendecomposition of the discretized shape Laplacians. For the sake of simplicity we assume Φ and Ψ to be weighted versions of the true eigenfunctions, allowing us to consider the standard scalar product in all equations instead of the manifold inner product. Note that in this case the truncated bases matrices meet $\Phi^\top \Phi = \mathbf{I}_k$, but $\Phi\Phi^\top \neq \mathbf{I}_m$ (due to truncation), and similarly for Ψ.

Throughout this paper we take the simplifying assumption that nodes can be matched exactly; in other words, we assume either the same tessellation, or a uniform point

distribution on the two shapes. This assumption allows us to represent the pointwise map T as a binary matrix $\mathbf{P} \in \{0,1\}^{n \times m}$. In the specific case in which it is known that the given shapes originate from two different deformations of the same template, then $n = m$ and matrix $\mathbf{P}$ is expected to be a permutation. Such a setting arises, for instance, in shape modeling and texture transfer applications [KS04]. In the general nearly-isometric setting with $n \neq m$ it is always possible to consider an equal number of points on the two shapes (obtained, for instance, by farthest point sampling), and seek for a permutation among the two reduced point sets [MDK*16]. In the sequel we will work with the general case $n \neq m$, which also models the setting in which one of the two shapes has missing parts.

In matrix notation, the expression for c_{ij} (2) is written as

$$\mathbf{C} = \Psi^\top \mathbf{P} \Phi, \tag{3}$$

where $\mathbf{C} = (c_{ij}) \in \mathbb{R}^{k \times k}$. Note that the matrix $\mathbf{C}$ is now a rank-k approximation of T_F. The problem of recovering the point-wise map from a given functional map, which is the objective of this work, can now be stated mathematically as finding the matrix $\mathbf{P}$ from the sole knowledge of $\mathbf{C}, \Phi, \Psi$.

Assumptions. In order to be as general as possible, we do *not* assume the matching process which generated the given functional map to be known. Additionally, we do not require any particular properties of $\mathbf{C}$ (*e.g.*, orthogonality), hence allowing to recover maps between shapes undergoing arbitrary deformations. Our main focus is on the adoption of the functional map representation for the purpose of shape matching. Within this context, our only requirement is that the matrix representation $\mathbf{C}$ is given w.r.t. known bases Φ, Ψ. These bases, in turn, may come from diverse optimization processes such as [PP93, KBB*13, NVT*14].

3.1. The inverse problem of point-to-point map recovery

Considering the problem of recovering $\mathbf{P}$ from a given $\mathbf{C}$ according to (3) as a (highly underdetermined) ill-posed inverse problem, the natural recovery problem to consider is

$$\mathbf{P}^* = \arg\min_{\mathbf{P} \in \{0,1\}^{n \times m}} D(\mathbf{C}, \Psi^\top \mathbf{P} \Phi) + \alpha J^P(\mathbf{P}) \tag{4}$$

$$\text{s.t. } \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \ \mathbf{P}\mathbf{1} \leq \mathbf{1}, \tag{5}$$

for a suitable measure of distance D, a regularization function J^P to possibly impose some kind of desired smoothness of the transformation, and a regularization parameter α determining a trade-off between fidelity and regularity. In general, the minimization of (4)-(5) can be very challenging. As we showed in [RMC15], for $n = m$ and by considering the equality $\mathbf{P}\mathbf{1} = \mathbf{1}$ we get a quadratic assignment problem as a special case. We start our discussion by analyzing the currently adopted approximations of the latter problem.

Mapping indicator functions. The simplest and most direct way for reconstructing the pointwise map T from T_F consists in mapping highly peaked Gaussian functions δ_x for each point $x \in \mathcal{M}$ via T_F, obtaining the image $g = T_F(\delta_x)$, and then declaring $T(x) \in \mathcal{N}$ to be the point at which g attains the maximum [OBCS*12] (such a procedure applies without changes to the partial case). Such a method, however, suffers from at least two drawbacks. First, it requires mapping delta functions for all shape points, which can get expensive for large meshes. Second, the low-pass filtering due to the basis truncation has a delocalizing effect on the maximum of g (see inset), affecting the quality of the final pointwise assignment.

Linear assignment problem. Instead of applying general greedy methods or relaxation techniques to (4), let us recall some observations from [OBCS*12]. In particular, note that the delta function $\delta_x : \mathcal{M} \to \{0,1\}$ around point x has coefficients $(\phi_i(x))_{i=1,\dots,k}$ in the Laplacian basis. This can be seen by observing that the inner product $\Phi^\top \delta_x$ is selecting the column of $\Phi^\top$ corresponding to point x. Therefore, the image via T_F of all indicator functions centered at points of $\mathcal{M}$ is simply given by $\mathbf{C}\Phi^\top$. Recovering the point-to-point map could then be solved by finding the correspondence between the columns of $\mathbf{C}\Phi^\top$ and the columns of $\Psi^\top$.

In the full-to-full case with $n = m$, measuring the proximity between these columns in the ℓ^2 sense gives rise to the linear assignment problem (LAP):

$$\min_{\mathbf{P} \in \{0,1\}^{n \times n}} \|\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 \tag{6}$$

$$\text{s.t. } \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \ \mathbf{P}\mathbf{1} = \mathbf{1}. \tag{7}$$

We refer to the Appendix for a proof of equivalence between (6)-(7) and the LAP in standard form.

Although the problem above admits a polynomial time solution [Kuh55], typical shape sizes (several thousands of vertices) make computing this solution very expensive in practice [BÇ99]. Furthermore, the LAP formulation requires an equal number of samples on the two shapes and the absence of partiality.

Nearest neighbors. The authors of [OBCS*12] circumvent the computational costs of the above approach by proposing a nearest-neighbor technique for the recovery of the point-to-point correspondence. In the light of our previous analysis their idea is to consider the matching of every point, *i.e.*, column of $\mathbf{C}\Phi^\top$, to its nearest neighbor in $\Psi^\top$ separately.

Mathematically, the nearest-neighbor approach can be seen as a relaxation of problem (6)-(7), in which one seeks

for the best *left-stochastic* approximation $\mathbf{P}$, *i.e.*,

$$\min_{\mathbf{P} \in \{0,1\}^{n \times m}} \|\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 \tag{8}$$

$$\text{s.t. } \mathbf{P}^\top \mathbf{1} = \mathbf{1}. \tag{9}$$

In other words, in comparison to (7) one omits the constraint of $\mathbf{P}\mathbf{1} = \mathbf{1}$. The omission allows to minimize the problem above by *independently* solving for columns of $\mathbf{P}$, one per query. Note that such a separable optimization approach may produce one-to-many mappings as a result of the recovery process. Moreover, it is an asymmetric method: looking for nearest neighbors from $\Psi^\top$ to $\mathbf{C}\Phi^\top$, or vice-versa, will in general produce different results. The nearest-neighbor approach was also applied in [RCB*16] for the partial case.

Our method. The analysis we provided above puts in evidence two major drawbacks, namely: 1) The linear assignment and the nearest-neighbor approaches rely on the assumption that the input functional map $\mathbf{C}$ aligns well the columns of $\Phi^\top$ with those of $\Psi^\top$ in the ℓ^2 sense. 2) None of the above approaches incorporates regularity assumptions for the alignment process, *i.e.*, the regularization term J^P in the general inverse problem formulation (4) was omitted.

We propose to cast the point-to-point map recovery as a probability density estimation problem to obtain both, a better measure of proximity than the ℓ^2 distance and a tool to impose regularity assumptions on the alignment map. Within our model, we interpret the columns of $\mathbf{C}\Phi^\top$ as modes of a *continuous* probability distribution defined over $\mathbb{R}^k$ (the embedded functional space), while columns of $\Psi^\top$ constitute the data, *i.e.*, a discrete sample drawn from the distribution. The task is then to align the modes to the data, such that the point-to-point mapping can be recovered as the maximum posterior probability.

As a model for the distribution we consider a Gaussian mixture (GMM) with m components, having the columns of $\mathbf{C}\Phi^\top$ as centroids in $\mathbb{R}^k$. For simplicity, we assume the components have uniform weight $\frac{1}{m}$, and equal covariances σ^2. Furthermore, in order to account for partiality, we relax the surjectivity requirement on the map and add a dummy node to the mixture. The purpose of this node is to absorb mismatches and outliers, acting as a slack variable; a similar approach was taken, *e.g.*, in [GR96, MS10, PBB*13]. With this model, the GMM density function is:

$$p(\mathbf{y}_j) = (1-w)\frac{1}{m}\sum_{i=1}^{m} p(\mathbf{y}_j|\mathbf{x}_i; \sigma^2) + w\frac{1}{m}, \tag{10}$$

where we write $\mathbf{y}_j$ and $\mathbf{x}_i$ to denote the columns of $\Psi^\top$ and $\mathbf{C}\Phi^\top$ respectively, and define $p(\mathbf{y}_j|\mathbf{x}_i; \sigma^2) = \frac{1}{(2\pi\sigma^2)^{k/2}}\exp(-\frac{\|\mathbf{y}_j - \mathbf{x}_i\|^2}{2\sigma^2})$. The coefficient $w \in [0,1]$ weighs the relative contribution of the GMM and the slack (the latter being modeled as a uniform distribution) and reflects our prior knowledge on the amount of overlap (*i.e.*, surface in common) between the two shapes.

Now let $R_\theta : \mathbb{R}^k \to \mathbb{R}^k$ denote the (unknown) transformation aligning the centroid locations $\mathbf{x}_i$ to the data points, according to a set of transformation parameters θ. The alignment problem can then be solved by maximizing the likelihood, or equivalently by minimizing the negative log-likelihood function

$$L(\theta, \sigma^2) = -\sum_{j=1}^{n} \log\left((1-w)\frac{1}{m}\sum_{i=1}^{m} p(\mathbf{y}_j|R_\theta(\mathbf{x}_i); \sigma^2) + w\frac{1}{m}\right). \tag{11}$$

Note that the argument that minimizes (11) can be also interpreted as the argument that minimizes the Kullback-Leibler (KL) divergence between a continuous GMM distribution (represented by the $\mathbf{x}_i = (\mathbf{C}\Phi^\top)_i$) and a mixture of Dirac distributions (represented by the $\mathbf{y}_j = (\Psi^\top)_j$). By following the notation of Eq. (4), the above probabilistic model can be written in matrix form as

$$\mathbf{P}^* = \underset{\sigma, \mathbf{P} \in [0,1]^{n \times m}, \mathbf{P}^\top \mathbf{1} = \mathbf{1}}{\arg\min} D_{\text{KL}}^\sigma(\mathbf{C}\Phi^\top, \Psi^\top \mathbf{P}) + \alpha J^P(\mathbf{P}). \tag{12}$$

The data term implements the loss (11), which measures the Kullback-Leibler divergence between the two point sets arising from the spectral representation of the two shapes. Before we detail our choice of the regularization $J^P(\mathbf{P})$ in the next paragraph let us mention that for a given solution $\mathbf{P}^*$ of (12), one could consider the posterior probabilities $(\mathbf{P}^*)_{ij} = p(\mathbf{x}_i|\mathbf{y}_j)$ in order to produce the final assignment. In our numerical experiments we, however, found the application of the nearest neighbor method to the non-rigidly aligned point clouds to yield superior results. Note that this results in a *binary* assignment $\mathbf{P}^*$, hence no further post-processing is needed.

Regularization. The above probabilistic model leaves some freedom for the specific choice of a transformation R_θ (see Eq. (11)). A simple example is choosing this transformation as a simple rotation, parametrized by θ. To be more flexible we instead propose to regularize the displacement field defined by $\mathbf{V} := \mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}$, which consists of the displacement vectors for aligning each point $(\mathbf{C}\Phi^\top)_i$ with its corresponding point $(\Psi^\top \mathbf{P})_i$. This corresponds to a general non-rigid transformation of one point set onto the other. Assuming that the true displacement field is locally smooth gives rise to using the Tikhonov regularization $\|\Gamma\mathbf{V}\|^2$ proposed in [YG89, MS10]. Here Γ denotes a low-pass operator promoting smoothly changing velocity vectors, and hence a locally coherent motion.

More precisely, it was shown in [MS10] that the minimization of $\|\Gamma\mathbf{V}\|^2$ is equivalent to regularizing

$$\sum_{l=0}^{\infty} w_l \int \|D^l v(x)\|^2 \, dx$$

for weights w_l and the derivative operator D such that

Figure 4: *A few iterations of the minimization process for problem* (13)*. The shape shown inside the plot is matched to the reference shape on the left. The curve shows the average geodesic error induced by the point-wise correspondence obtained at each iteration. The error is visualized as a heat map encoding distance to ground truth, growing from white to hot colors (white means zero error).*

$D^{2l}v = \Delta^l v$ and $D^{2l+1}v = \nabla\Delta^l v$, for Δ denoting the Laplace operator. Note that our regularization includes the penalization of $\|\nabla v(x)\|^2$, which is also known as the Dirichlet energy. Although the Dirichlet energy has been considered in the context of shape matching in [KBBV15] before, here we take a significantly different approach: While the Dirichlet energy was used in [KBBV15] to induce smooth images of indicator functions under the functional mapping, we use it as a *regularization* for the displacement field aligning the two point clouds arising from the spectral representation of the two shapes.

To illustrate this idea, the inset shows a smooth velocity field with coherent correspondences (blue) and a mismatch produced by simple nearest-neighbors (red). Combining the motion coherence regularization with the objective function we obtain

$$\mathbf{P}^* = \underset{\sigma, \mathbf{P} \in [0,1]^{n \times m}, \mathbf{P}^\top \mathbf{1} = \mathbf{1}}{\arg\min} D_{\mathrm{KL}}^{\sigma}(\mathbf{C}\Phi^\top, \Psi^\top \mathbf{P})$$
$$+ \lambda \|\mathbf{\Gamma}(\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P})\|^2, \quad (13)$$

where $\lambda > 0$ is a trade-off parameter between likelihood and regularity (set to $\lambda = 3$ in our experiments). We refer to Fig. 5 for a sensitivity analysis with respect to this parameter.

General alignment problems like (13) have been widely researched in computer vision, and several robust algorithms exist for these tasks [CR00, TK04, MS10, JV11]. Most of these approaches follow an iterative scheme, optimizing w.r.t. $\{\mathbf{V} \approx \mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}, \sigma\}$ and $\mathbf{P} = (p(\mathbf{x}_i|\mathbf{y}_j))_{ij}$ in an alternating fashion until convergence (EM algorithm [DLR77]). In our experiments, we used publicly available code from [MS10], which allows to optimize over smooth displacements as in Eq. (13).

In Fig. 4 we show a few iterations of the EM recovery process applied to a pair of nearly-isometric shapes, starting from a functional map obtained as described in Section 5.
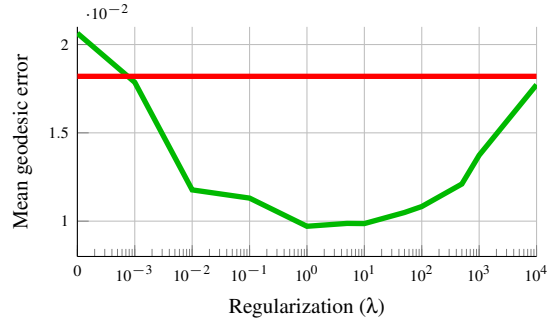

Figure 5: *Effect of the regularization parameter λ (see Eq.* (13)*) on the quality of the recovered point-wise map. For $\lambda \to \infty$ we enforce the Dirichlet energy of the displacement to be zero and hence restrict ourselves to a rigid alignment, which naturally leads to approximating the accuracy of ICP (red line). For $\lambda = 0$ (no regularization) we do not enforce consistent motion at all and obtain slightly worse results than ICP.*

4. Point-wise map refinement

As a general representation for shape correspondences, functional maps can be adopted to compactly encode (via Eq. (3)) dense point-to-point maps obtained by any matching algorithm such as [KLF11, RBW*14]. Therefore, one can consider improving an input noisy point-to-point map (i.e., to remove mismatches that are manifested as high-frequency noise) by following a simple two-step procedure: First, the input map is represented as a functional correspondence $\mathbf{C}$; then, the point-to-point map $\mathbf{P}$ is recovered again by following the robust recovery method we described in the previous section. Naturally, one can repeat such a strategy and iterate between updating the point-to-point correspondence and the functional map. Images of delta functions under the resulting 'refined' functional map will be sharper and exhibit less spread than the initial map.

Such an iterative procedure was first considered in [OBCS*12], where the authors proposed to use the classical Iterative Closest Point (ICP) algorithm [BM92]. The latter updates $\mathbf{P}$ according to the nearest-neighbor approach (8), followed by a refinement of $\mathbf{C}$ via

$$\min_{\mathbf{C} \in \mathbb{R}^{k \times k}} \|\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 \quad (14)$$

$$\mathbf{C}^\top \mathbf{C} = \mathbf{I}_k, \quad (15)$$

which is an orthogonal Procrustes problem. Intuitively, this can be seen as a rigid alignment in $\mathbb{R}^k$ between point sets $\Phi^\top$ and $\Psi^\top \mathbf{P}$ (see Fig. 7(b) for an example). The alternating minimization w.r.t. $\mathbf{C}$ and $\mathbf{P}$ is repeated until convergence.

Although the ICP approach described above allows to achieve significant improvements in terms of map accuracy, the orthogonality constraints (15) imposed on the functional

RC (35.85%)   **R**   **C** (15.98%)

Figure 6: *Solving problem (16) gives rise to the refined functional map shown on the left, where $k = 20$. The refined map attains a higher percentage of exact matches (reported in parentheses) than the initial map. The optimal transformation (middle) essentially delineates the refinement process as a perturbation of the identity.*

correspondence limit its applicability to a specific class of transformations, namely the class of *volume-preserving isometries* (see [OBCS*12] Theorem 5.1). Therefore, the method cannot be applied to refine maps between shapes undergoing arbitrary deformations.

For a more general refinement procedure, we drop the orthogonality constraints (15) and instead optimize over a general transformation **R** that is applied to the given **C** . Specifically, we consider the general problem

$$\min_{\mathbf{R} \in \mathbb{R}^{k \times k}} \|\mathbf{R}\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 + \beta J^C(\mathbf{R}), \qquad (16)$$

for a regularization functional J^C which could encourage **RC** to correspond to a smooth transformation, or could require **R** to be a small perturbation of the identity. While for the specific example

$$J^C(\mathbf{R}) = \begin{cases} 0 & \text{if } (\mathbf{R}\mathbf{C})^\top \mathbf{R}\mathbf{C} = \mathbf{I}_k, \\ \infty & \text{else,} \end{cases}$$

problem (16) coincides with the rigid alignment problem arising from the constraint (15), a less restrictive choice for the regularization functional J^C makes the method suitable for recovering functional maps for non-isometric shape matching problems.

In our experiments we found that when (16) is combined with our proposed regularized probabilistic model (13) (which provides the input assignment **P**), it is sufficient to simply update **R** in a least-squares sense: In other words, we simply set $\beta = 0$ and solve the resulting problem. By doing so, **R** is determined to be a perturbation of the identity as illustrated in Fig. 6. The fact that **C** is refined in a *non-rigid* fashion can improve the refinement results significantly in comparison to orthogonal updates of **C** as illustrated in 3d in Fig. 7. A similar alignment procedure was considered in [MHK*08, LGPC13] in lower dimensions for articulated object matching.

Partial matching. In the partial setting, the refinement process is complicated by the fact that the two shapes are related by a non-rigid transformation mapping the spectral em-


reference   (a)   (b)   (c)

Figure 7: *The refinement of a rank-k functional map can be seen as the alignment of two point sets in $\mathbb{R}^k$. Here the reference shape (left) is matched to the pose shown in the three columns. In the top row we illustrate the action of different refinement methods when $k = 3$, with the two point sets being plotted as red and white point clouds. (a) Initial map; (b) Solution after orthogonal refinement (14); (c) Solution obtained with our approach (16). In the bottom row we show the color-coded pointwise matches w.r.t. the reference shape (corresponding points have the same color) when $k = 50$.*

bedding of the partial shape ($\Psi^\top$) to some *unknown subset* of the embedding of the full shape ($\mathbf{C}\Phi^\top$). The two spectral embeddings overlap only partially, with a pre-alignment provided by the initial **C** (see Fig. 8 left). In [RCB*16]


Figure 8: Left*: When the two shapes have missing parts, their embeddings in $\mathbb{R}^k$ overlap partially. Shown here are the first three dimensions of $\Psi^\top$ (red) and $\mathbf{C}\Phi^\top$ (white) for $\mathcal{N}_{\text{part}}$ and $\mathcal{M}$ from Fig. 2. Right: Multiplication by $\mathbf{C}$ in the partial case jointly aligns and projects the spectral embedding of the full shape onto a lower-dimensional subspace.*

the authors showed how to compute such a pre-alignment, and subsequently used the nearest-neighbor approach to recover a partial pointwise map. Specifically, their optimal **C** is such that $\mathbf{C}\mathbf{C}^\top \approx \begin{pmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. This property induces a slanted-diagonal structure on **C**, with the top non-zero block $\mathbf{C}_r$ being *orthogonal* (see Fig. 2) – hence resembling the ICP procedure of (14)-(15).

We note here that multiplying $\Phi^\top$ by a slanted-diagonal **C** has a double action. First, as in the full case, the application of **C** has the effect of aligning the point sets $\Phi^\top$ and $\Psi^\top$;

by the requirement that $\mathbf{C}_r\mathbf{C}_r^\top \approx \mathbf{I}$, this alignment is asked to be *as rigid as possible*. Second, the zero block of $\mathbf{C}$ sets the last $k-r$ coordinates of the aligned point set $\mathbf{C}\Phi^\top$ to zero. This corresponds to an orthogonal projection onto a r-dimensional subspace of $\mathbb{R}^k$, where the point sets align well. This joint effect is illustrated for $k=3$ and $r=2$ in Fig. 8.

A similar interpretation can be given to our probabilistic recovery procedure. We now seek for the best *non-rigid* alignment between $\Psi^\top$ and the *projected* point set $\mathbf{C}\Phi^\top$. Following Eq. (10), we also need an estimate $w \in [0,1]$ for the overlap between full and partial shape. While this can be easily obtained by computing the area ratio of the two shapes, a better option is to use once again an estimate of the rank of $\mathbf{C}$ (in fact, as a consequence of Weyl's law, we get $|\mathcal{N}| \approx \frac{r}{k}|\mathcal{M}|$ [RCB*16]). Our choice is motivated by the fact that the ratio of r over k is more aware of the deformation relating the two shapes than the simple ratio of surface areas. Hence, we set $w = 1 - \frac{r}{k}$. Note that this estimate is general in that it also applies to the full-to-full case, since in this case one obtains $w \approx 0$.

5. Experimental evaluation

We compare our iteratively refined probabilistic point-wise map recovery method with the iterative refinement procedure of Ovsjanikov *et al.* [OBCS*12] (denoted as ICP, see Section 4), and with the ICP-like procedure described in [RCB*16] for partial matching. The two methods are, to the best of our knowledge, the only existing alternatives to date. All algorithms were implemented in Matlab/MEX and executed (single-core) on an Intel i7-3770 3.4GHz cpu with 32GB memory. We also realized an accelerated gpu implementation of our method, and tested it on an Nvidia GTX 750s. Code and data for our method are available for download at http://sites.google.com/site/erodola/publications/.

As a measure of error, we use the quantitative criterion that was introduced in [KLF11] to evaluate the quality of point-wise maps. The input quantity in our case is a functional map $\mathbf{C}$, which is then converted to its point-wise counterpart using each method. We plot cumulative curves showing the percent of matches which have normalized geodesic error smaller than a variable threshold.

Each method is evaluated quantitatively on the FAUST [BRLB14] and KIDS [RBW*14] datasets, and qualitatively on the TOSCA [BBK08] dataset. The shapes from TOSCA and KIDS were remeshed independently to $\sim$10K vertices each by iterative edge contractions [GH97], so as to avoid identical meshings and make the data more challenging. The three datasets include isometric as well as non-isometric shapes; in particular, FAUST and KIDS also include point-to-point ground truth matches between shapes belonging to different classes. For the evaluation on partial shapes we used the TOSCA "cuts" dataset that was introduced in [RCB*16].


Figure 9: Left: *Comparisons with nearly-isometric shapes; in this setting, orthogonal refinement (ICP) already provides accurate point-to-point recovery, which can be further improved by applying our algorithm.* Right: *Reference for normalized geodesic error. Colors are labeled by distance according to the legend.*


Figure 10: *Comparisons with non-isometric shapes. In this case, the lack of orthogonality in $\mathbf{C}$ drives ICP to poor solutions. This is especially evident with the KIDS dataset, where the shapes have a larger variety in the geometry; here, our method improves upon ICP by 40%. Note the different error ranges for the two plots.*

Comparisons. The functional maps used in the comparisons are constructed by solving a least-squares system $\mathbf{C}\mathbf{A} = \mathbf{B}$, where matrices $\mathbf{A}$ and $\mathbf{B}$ contain the Fourier coefficients of indicator functions for corresponding points on the two shapes. The point-to-point correspondence is established using the ground truth, while the indicator functions are approximated as highly peaked Gaussians around each point. We use 35 sparse matches for constructing the initial $\mathbf{C}$, distributed according to a farthest point sampling using the Euclidean metric. This way, we simulate a matching process that provides reasonably good solutions for further refinement. Note that this process is more realistic than what we did in [RMC15], where the functional map was initialized by matching corresponding regions [RBC14].

We show comparisons both in the intra-class (near-

Figure 11: Left*: Performance of each method with functional maps constructed from k ground-truth matches, as the value of k increases from 3 to 35.* Right*: Performance at increasing outlier ratio; in this experiment, random pairs of points are injected in the set of matches used to construct the input functional maps. Observe how both our method and ICP remain stable at increasing noise levels.*

Figure 12: *Comparisons on the FAUST dataset, where each problem is initialized with a functional map obtained with the random forest method of [RBW*14].*

isometric) and inter-class (non-isometric) settings. In the former case, we use the complete intra-class template subset of FAUST, which consists of 90 pairs of humans in different poses subdivided into 10 different individuals. All shapes have $n = 6890$ points, and the functional map is computed with $k = 30$ basis functions. Since we only use 35 known matches to construct the initial $\mathbf{C}$, this setup makes it challenging for any method to recover the exact point-wise map. Results are reported in Fig. 9 (left), where we also included plain nearest-neighbors (NN) as a baseline. From the plotted curves we can see that orthogonal ICP is performing slightly better than our method on near-isometric deformations, since the approach is specifically tailored for this particular case. However, initializing our method with the output of ICP allows to achieve 10% further improvement on average.

In Fig. 10 we show the same curves for the non-isometric case (inter-class matching) on the FAUST and KIDS datasets, consisting respectively of 90 and 45 pairs of different individuals under different poses. In this case, the accuracy of orthogonal refinement worsens due to the different properties of the input functional maps, which now relate shapes under *non* volume-preserving transformations. On FAUST data, the discrepancy with our method increases to 20%, while it rises to 40% with KIDS shapes.

We additionally studied the change in performance of each method when fed with noisy functional maps as input (Figure 11, see caption for details). As we can see from the plots, the relative performance of each method is not strongly affected by the quality of the input. Finally, in Fig. 12 we plot the results obtained by each method on the FAUST dataset, when the initial functional map comes from an actual matching pipeline (namely the random forest method of [RBW*14]). We see that the accuracy improvements reflect the ones observed in the simulated set-
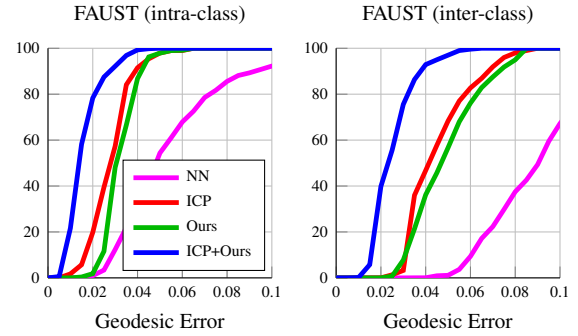
ting, with even a bigger improvement over ICP (>30%) in the inter-class experiments. Additional qualitative comparisons are shown in Fig. 13.

Rank. In a second set of experiments, we evaluate the capability of each method to recover point-wise maps from functional maps of increasing rank. In this setting, we assume the input functional map to be as accurate as possible, and for this purpose we construct it explicitly as $\mathbf{C} = \Psi^\top \mathbf{P} \Phi$, where $\mathbf{P}$ is the ground-truth permutation among the vertices of the two shapes. We do so for a pair of approximately isometric shapes, so that the respective eigenbases Φ and Ψ are as compatible as possible, and further orthogonal refinement is not needed (indeed, applying ICP in this setting actually yielded worse results in our tests).

The results are shown in Fig. 14. As the number k of basis functions used on the two shapes (*i.e.*, the rank of $\mathbf{C}$) increases, so does the amount of exact correspondences recovered by each method. This is also true for the simple approach mapping indicator functions (Max), since the smoothing effect due to basis truncation is reduced at increasing values of k.

Our method allows to recover up to 20% more *exact* matches than the nearest-neighbors approach. In particular, with $k = 100$ (a commonly used value in most shape matching pipelines) we are able to perfectly reconstruct 75% of the rows/columns of $\mathbf{P}$ (a 6890×6890 matrix in this example).

Partial matching. In Fig. 15 we show comparisons with the ICP-like approach of Rodolà *et al.* [RCB*16]. The comparisons are performed on the *cat* and *dog* classes of the TOSCA "cuts" dataset, consisting of 95 matching problems among near-isometric shapes with missing parts. Note that, due to the nearly isometric transformations, ICP already gets accurate results which we further improve upon by 10% on average. This is consistent with what we observe in the near-isometric case among full shapes (see Fig. 9).

Figure 13: *Refinement examples in different matching scenarios. In each row we show the source shape (left), followed by the map errors produced by ICP (middle) and our method (right); the error is visualized as a heat map encoding distance to the ground truth, growing from white to black. The two methods perform comparably well in the near-isometric case, but the orthogonal refinement of ICP yields large errors with more general deformations and missing parts.*

Complexity issues. The time performance of our method depends on two factors: the number of shape points n, and the size of the functional map k. As we also show in Fig. 4, typically a few iterations of the EM algorithm are sufficient to reach accurate solutions, and in practice we used 5 iterations in all our experiments. In the common case where $n = 10,000$ and $k = 30$, the cpu implementation of our

n	cpu	gpu (double)	gpu (single)
7K	30.81	10.70	**6.38**
10K	92.02	21.80	**12.85**

Table 1: *Average running times (in seconds) for 5 iterations of our algorithm, on shapes from the FAUST (first row) and TOSCA (second row) datasets.*


Figure 14: *Percentage of* exact *correspondences (solid curves) recovered from ground-truth functional maps of increasing rank among the two shapes shown on the right. We also report the percent of correspondences with geodesic error smaller than* 0.02 *(dashed curves).*

method takes on average 1 min. 30 sec. to converge, while ICP using exact search structures adds up to ∼15 sec.

In order to improve its efficiency, we also implemented our algorithm in gpu using cuBLAS [cub15], CULA [cul14], and matrix slicing to circumvent memory limitations. Runtime comparisons with the cpu implementation on two datasets of different sizes are reported in Tab. 1. In the table we include the runtimes obtained with double- and single-precision arithmetics, where the latter is observed to improve runtime performance at a negligible decrease in accuracy.

6. Discussion and conclusions

In this paper we formulated a general variational recovery approach for the inverse problem of computing point-to-point correspondences from a given functional map. We introduced a probabilistic model for point-wise map recovery and considered a refinement of the functional map that does not rely on the assumption of isometric shapes. In this context, the refinement procedure can be seen as a non-rigid alignment between the spectral embeddings of the two shapes. We further showed how this method can be naturally adapted to a setting in which one of the two shapes has missing parts. The experimental results showed that the proposed approach yields up to 10% accuracy improvements under partiality transformations, up to 40% with non-isometric deformations, and reaching up to 75% *exact* point-to-point matches under good initializations.

The main limitation of our recovery method lies in the fact that – similarly to previous approaches – the optimization procedure is biased towards one of the two shapes, as one can for instance see from the interpretation of minimizing the (non-symmetric) Kullback-Leibler divergence. Incorporating a symmetric term in the probabilistic model to remove this bias (i.e., solving for the map and its inverse simultaneously) represents a possibility. Second, the partial similarity
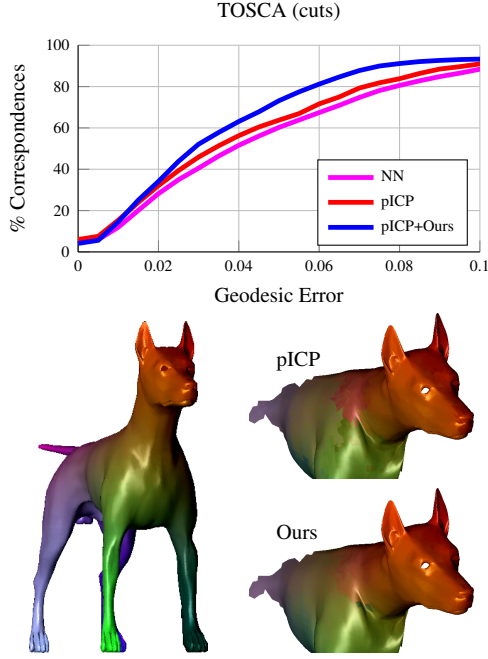
Figure 15: *Comparisons with partial shapes on a subset of the TOSCA "cuts" dataset. pICP denotes the ICP-like recovery technique of [RCB*16] for part-to-whole matching. In the bottom row, corresponding points have the same color.*

setting considered in this paper allows only one of the two shapes to have missing parts. Deformable matching of partial shapes in presence of occlusions and clutter remains a challenging problem, and an interesting direction of further research.

Appendix

We prove that the following problem:

$$\min_{\mathbf{P}\in\{0,1\}^{n\times n}} \|\mathbf{C}\Phi^\top - \Psi^\top\mathbf{P}\|_F^2 \tag{17}$$

$$\text{s.t.}\ \ \mathbf{P}^\top\mathbf{1} = \mathbf{1},\ \mathbf{P}\mathbf{1} = \mathbf{1} \tag{18}$$

can be rewritten equivalently as a linear assignment problem in standard form. First, note that the constraints on $\mathbf{P}$ require this matrix to be a permutation. Next, observe that

$$\|\mathbf{C}\Phi^\top - \Psi^\top\mathbf{P}\|_F^2 = \|\mathbf{C}\Phi^\top\|_F^2 + \|\Psi^\top\mathbf{P}\|_F^2 - 2\langle\mathbf{C}\Phi^\top, \Psi^\top\mathbf{P}\rangle$$
$$= \|\mathbf{C}\Phi^\top\|_F^2 + \|\Psi^\top\|_F^2 - 2\langle\Psi\mathbf{C}\Phi^\top, \mathbf{P}\rangle$$

holds for permutation matrices $\mathbf{P}$. Since we are interested only in the argument that minimizes (17) subject to (18), we can discard the first two summands and thus arrive at the linear assignment problem:

$$\min_{\mathbf{P}\in\{0,1\}^{n\times n}} -\langle\Psi\mathbf{C}\Phi^\top, \mathbf{P}\rangle \tag{19}$$

$$\text{s.t.}\ \ \mathbf{P}^\top\mathbf{1} = \mathbf{1},\ \mathbf{P}\mathbf{1} = \mathbf{1}. \tag{20}$$

References

[BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 8

[BÇ99] BURKARD R. E., ÇELA E.: *Linear Assignment Problems and Extensions*. Springer US, Boston, MA, 1999, pp. 75–149. doi:10.1007/978-1-4757-3023-4_2. 4

[BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *TPAMI 14*, 2 (1992), 239–256. 6

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (2014), IEEE, pp. 3794–3801. 8

[CR00] CHUI H., RANGARAJAN A.: A feature registration framework using mixture models. In *Proc. MMBIA* (2000), IEEE, pp. 190–197. 6

[CRB*16] COSMO L., RODOLÀ E., BRONSTEIN M. M., ET AL.: SHREC'16: Partial matching of deformable shapes. In *Proc. 3DOR* (2016). 3

[CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSELLO A., BRONSTEIN M. M.: Matching deformable objects in clutter. In *Proc. 3DV* (Oct 2016), pp. 1–10. 2

[cub15] cuBLAS library user guide v7.5, 2015. URL: http://docs.nvidia.com/cuda/cublas/index.html. 10

[cul14] CULA reference manual, 2014. URL: http://www.culatools.com. 10

[DLR77] DEMPSTER A. P., LAIRD N. M., RUBIN D. B.: Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Series B 39*, 1 (1977), 1–38. 6

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216. 8

[GR96] GOLD S., RANGARAJAN A.: A graduated assignment algorithm for graph matching. *TPAMI 18*, 4 (1996), 377–388. 5

[GT15] GASPARETTO A., TORSELLO A.: A statistical model of Riemannian metric variation for deformable shape analysis. In *Proc. CVPR* (2015), IEEE, pp. 1219–1228. 2

[HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph. 33*, 4 (July 2014), 36:1–36:11. 2

[JV11] JIAN B., VEMURI B. C.: Robust point set registration using Gaussian mixture models. *TPAMI 33*, 8 (2011), 1633–1645. 6

[KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Comput. Graph. Forum 32*, 2pt4 (2013), 439–448. 2, 4

[KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proc. CVPR* (2015), IEEE, pp. 905–914. 2, 3, 6

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG 30*, 4 (2011), 79. 6, 8

[KS04] KRAEVOY V., SHEFFER A.: Cross-parameterization and compatible remeshing of 3d models. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 861–869. doi:10.1145/1015706.1015811. 4

[Kuh55] KUHN H. W.: The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly 2*, 1–2 (March 1955), 83–97. 4

[KZHCO11] KAICK O. V., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Comput. Graph. Forum 30*, 6 (2011), 1681–1707. 1

[LGPC13] LOMBAERT H., GRADY L., POLIMENI J. R., CHERIET F.: Focusr: Feature oriented correspondence using spectral regularization–a method for precise surface matching. *TPAMI 35*, 9 (Sept. 2013), 2143–2160. 7

[LRB∗16] 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

[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. URL: http://doi.acm.org/10.1145/2897824.2925913, doi:10.1145/2897824.2925913. 4

[MHK∗08] MATEUS D., HORAUD R., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR* (June 2008), IEEE, pp. 1–8. 7

[MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *TPAMI 32*, 12 (2010), 2262–2275. 5, 6

[NVT∗14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. *Comput. Graph. Forum 33*, 5 (2014), 1–10. 4

[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 (2012), 30:1–30:11. 1, 2, 3, 4, 6, 7, 8

[PBB∗13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Comput. Graph. Forum 32*, 2pt4 (2013), 459–468. 2, 3, 5

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental mathematics 2*, 1 (1993), 15–36. 4

[RBC14] RODOLÀ E., BULÒ S. R., CREMERS D.: Robust region detection via consensus segmentation of deformable shapes. *Comput. Graph. Forum 33*, 5 (2014), 97–106. 8

[RBW∗14] RODOLÀ E., BULÒ S. R., WINDHEUSER T., VESTNER M., CREMERS D.: Dense non-rigid shape correspondence using random forests. In *Proc. CVPR* (2014), IEEE, pp. 4177–4184. 2, 6, 8, 9

[RCB∗16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Comput. Graph. Forum* (2016). 2, 3, 5, 7, 8, 9, 11

[RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Proc. VMV* (2015), Bommes D., Ritschel T., Schultz T., (Eds.), The Eurographics Association, pp. 25–32. 2, 4, 8

[ROA∗13] RUSTAMOV R. M., OVSJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph. 32*, 4 (2013), 72:1–72:12. 2

[TK04] TSIN Y., KANADE T.: A correlation-based approach to robust point set registration. In *Proc. ECCV*, vol. 3023. Springer Berlin Heidelberg, 2004, pp. 558–569. 6

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

[WHG13] WANG F., HUANG Q., GUIBAS L.: Image co-segmentation via consistent functional maps. In *Proc. ICCV* (2013), IEEE, pp. 849–856. 2

[YG89] YUILLE A. L., GRZYWACZ N. M.: A mathematical analysis of the motion coherence theory. *Int. J. of Comp. Vision 3*, 2 (1989), 155–175. 5