

Point-wise Map Recovery and Refinement from Functional Correspondence

E. Rodolà, M. Moeller, D. Cremers

TU Munich, Germany

Abstract

Since their introduction in the shape analysis community, functional maps have met with considerable success due to their ability to compactly represent dense correspondences between deformable shapes. Despite the numerous advantages of such representation, however, the problem of converting a given functional map back to a point-to-point map has received a surprisingly limited interest. In this paper we analyze the general problem of point-wise map recovery from arbitrary functional maps. In doing so, we rule out many of the assumptions required by the currently established approach – most notably, the limiting requirement of the input shapes being nearly-isometric. We devise an efficient recovery process based on a simple probabilistic model. Experiments confirm that this approach achieves remarkable accuracy improvements in very challenging cases.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Computer Graphics]: Scene Analysis—Shape

1. Introduction

Shape matching is a widely researched topic in computer vision and graphics, and a diverse range of techniques that tackle the problem of correspondence have been proposed during the years [VKZHC01]. Of particular interest is the case in which the input shapes are allowed to undergo non-rigid deformations, which are typically assumed to be approximately isometric. Recent advancements in this area include the seminal work of Ovsjanikov *et al.* [OBCS*12], who proposed modeling *functional correspondence* between shapes; in this view, the focus shifts from studying point-wise mappings to the definition of a linear operator (the *functional map*) relating spaces of functions on the two shapes. The classical point-wise representation constitutes then a special case in which the functional map corresponds delta-functions to delta-functions. A major advantage of the functional representation lies in the linearity of the operator: the functional map admits a matrix representation which can be made *compact* under a proper choice of bases for the two functional spaces. In [OBCS*12] the authors advocated the use of the Laplacian eigenfunctions as the natural basis for smooth functions on the respective shapes; with this choice, one is allowed to “truncate” the representation by using the first few basis functions and still obtain a good approximation to the underlying point-wise correspondence.




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.

Follow-up work by several authors showed how to extend the framework to deal with non-isometric deforma-

tions [PBB*13, KBB*13, RRBW*14, KBBV15], partial similarity [RCB*15], 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 established approach, originally proposed in [OBSC*12], operates by formulating the conversion problem as a nearest-neighbor search in the embedded functional space; as we show in the following sections the approach works well if the functional map is accurate enough, with significant decrease in accuracy as the number of basis functions is reduced. The resulting point-to-point map can be iteratively refined by following a simple procedure, but this can only be done under specific assumptions on the initial functional map. A similar refinement technique was recently applied for near-isometric partial matching [RCB*15], and in a correspondence-less setting for shape retrieval tasks [GT15]. Finally, the point-wise recovery problem was sidestepped in [KBBV15] by adopting a soft error criterion to evaluate the quality of functional maps *without* converting them to a point-wise counterpart.

Contribution. In this paper, we consider the problem of accurate point-wise map recovery from a given functional map. The key contributions can be summarized as follows:

- We provide the first rigorous analysis of the point-wise map recovery problem. In particular, we show how a simple modification to the baseline approach can lead to consistent improvements.
- We introduce a simple probabilistic model for map recovery *and* refinement. Our model does not impose any assumption on the input functional map, as well as on the specific choice of a functional basis on the two shapes.
- Our method is efficient, and significantly outperforms the existing method in both, the nearly-isometric and the inter-class settings.

2. Background

We model shapes as compact connected two-dimensional Riemannian manifolds \mathcal{M} (possibly with boundary) endowed with the standard measure μ induced by the volume form. Let $L^2(\mathcal{M}) = \{f : \mathcal{M} \rightarrow \mathbb{R} \mid \int_{\mathcal{M}} f^2 d\mu < \infty\}$ denote the space of square-integrable functions on \mathcal{M} , and let $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} fg d\mu$ be the standard manifold inner product. The space (\mathcal{M}, μ) features the symmetric Laplace-Beltrami operator (or Laplacian) $\Delta_{\mathcal{M}} : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{M})$, which provides us with all the tools of Fourier analysis on our manifold. In particular, this operator admits 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 on $L^2(\mathcal{M})$.

Drawing an analogy with classical signal processing theory, the eigenfunctions are often referred to as *manifold harmonics*, and the associated eigenvalues assume the interpretation of frequencies [Tau95]. Any function $f \in L^2(\mathcal{M})$ then admits a Fourier series expansion as

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

Functional correspondence. Let us be given two manifolds \mathcal{M} and \mathcal{N} , and let $T : \mathcal{M} \rightarrow \mathcal{N}$ be a bijective mapping between them. Departing from the traditional point-centric setting, Ovsjanikov et al. [OBSC*12] introduced the notion of functional map between two shapes as the linear operator $T_F : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{N})$, mapping functions on \mathcal{M} to functions on \mathcal{N} via the composition $T_F(f) = f \circ T^{-1}$. The approach is a natural generalization to classical point-wise correspondence, which can be seen as the special case in which T_F maps indicator functions to indicator functions.

Let $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$ denote orthonormal bases on $L^2(\mathcal{M})$ and $L^2(\mathcal{N})$ respectively. The functional correspondence with respect to these bases can be expressed as follows, for some function $f \in L^2(\mathcal{M})$:

$$\begin{aligned} T_F(f) &= T_F \left(\sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i \right) = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} T_F(\phi_i) \\ &= \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \underbrace{\langle T_F(\phi_i), \psi_j \rangle_{\mathcal{N}}}_{c_{ij}} \psi_j. \end{aligned} \quad (2)$$

Thus, the action of T_F amounts to linearly transforming the expansion coefficients of f from basis $\{\phi_i\}_{i \geq 1}$ onto basis $\{\psi_j\}_{j \geq 1}$. The transformation is encoded in the coefficients c_{ij} , providing a representation for T_F as the (possibly infinite) matrix $\mathbf{C} = (c_{ij})$. Seeking a functional correspondence among the two shapes then amounts to solving for the unknown \mathbf{C} that better preserves certain mapping constraints [OBSC*12] or manifesting regularity at different levels [RRBW*14, KBBV15].

Basis truncation. A natural choice for a basis on the two shapes is given by the eigenfunctions $\{\phi_i\}_{i \geq 1}$, $\{\psi_j\}_{j \geq 1}$ of the respective Laplacians (harmonic basis). The basis functions are said to be compatible if the equality $\psi_j = \pm \phi_i \circ T^{-1}$ holds (approximately) for all $i \geq 1$, which is the case with the manifold harmonics when the shapes are related by a near-isometry. If the deformation is far from isometric, compatible basis functions can still be computed ad-hoc for the two manifolds, based on a minimal set of coupling functions (e.g., a sparse set of point-to-point matches) [KBB*13].

Assuming to have stable and compact (namely, harmonic) bases at disposal, in [OBSC*12] the authors proposed to truncate the series (2) at the first k coefficients, resulting in a $k \times k$ matrix \mathbf{C} approximating the full map in a compact way. The reduced representation greatly simplifies correspondence-based tasks (e.g., shape matching); 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 bijection T from its functional representation T_F .

3. Point-wise map recovery

Let us now consider the discretized problem and assume that \mathcal{M} and \mathcal{N} are represented by a triangular mesh with n nodes each. Let matrices $\Phi, \Psi \in \mathbb{R}^{n \times k}$ contain the first k basis functions for the two shapes respectively, represented as column vectors. Note that, due to truncation, we have that $\Phi^\top \Phi = \mathbf{I}_k$, but $\Phi \Phi^\top \neq \mathbf{I}_n$, and similarly for Ψ . We assume that the bijection $T : \mathcal{M} \rightarrow \mathcal{N}$ is known. For the sake of simplicity, we additionally assume that T can be represented as a permutation matrix $\mathbf{P} \in \{0, 1\}^{n \times n}$, which means that the correspondence originates from two different deformations of the same template. In this case the expression for c_{ij} in (2) can be equivalently rewritten as:

$$\mathbf{C} = \Psi^\top \mathbf{P} \Phi, \quad (3)$$

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 objective of any process converting the functional map back into a point-wise map representation, which is the focus of this work, is to recover the permutation \mathbf{P} from the sole knowledge of \mathbf{C} , Φ , and Ψ .

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 from \mathbf{C} (e.g., orthogonality), hence allowing to recover maps between shapes undergoing arbitrary deformations. 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].

Mapping indicator functions.

The simplest and most direct way for reconstructing the bijection T from the associated functional map T_F consists in mapping indicator functions $\delta_x : \mathcal{M} \rightarrow \{0, 1\}$ 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 method, however, suffers from at least two drawbacks. First, it requires constructing and mapping indicator functions for all shape points, which can get easily 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 figure), negatively affecting the quality of the final correspondence.



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

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

$$\text{s.t. } \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \mathbf{P} \mathbf{1} = \mathbf{1}, \quad (5)$$

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. Note that the constraints in (5) might be adapted to account for local area elements as in [ADK13].

Unfortunately, the minimization of (4) will be extremely challenging in general. Consider for instance the case of no regularization, $J^P \equiv 0$, a zero functional map $\mathbf{C} = 0$, and the distance measure being the squared Frobenius norm. In this case the minimization problem becomes

$$\begin{aligned} \min_{\substack{\mathbf{P}_{i,j} \in \{0, 1\} \\ \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \mathbf{P} \mathbf{1} = \mathbf{1}}} \|\Psi^\top \mathbf{P} \Phi\|_F^2 &= \min_{\substack{\mathbf{P}_{i,j} \in \{0, 1\} \\ \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \mathbf{P} \mathbf{1} = \mathbf{1}}} \|\mathbf{Q} \vec{\mathbf{P}}\|_2^2, \\ &= \min_{\substack{\mathbf{P}_{i,j} \in \{0, 1\} \\ \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \mathbf{P} \mathbf{1} = \mathbf{1}}} \vec{\mathbf{P}}^\top \mathbf{Q}^\top \mathbf{Q} \vec{\mathbf{P}}, \end{aligned}$$

for the right hand sides denoting vectorized equations with $\mathbf{Q} = \Phi^\top \otimes \Psi$ being formed by the Kronecker product between Φ^\top and Ψ . For arbitrary \mathbf{Q} the above problem is a particular reformulation of the quadratic integer programming problem (see Equation (8) in [BcPP98]). Since the latter was shown to be NP-hard we cannot expect to find exact solutions to the above problem with reasonable computational complexity in general. Therefore, we turn our attention to approximations of (4).

Linear assignment problem. Instead of applying general greedy methods or relaxation techniques to (4), let us recall some observations from [OBCS*12] regarding the general structure of (3): In particular, note that the indicator function $\delta_x : \mathcal{M} \rightarrow \{0, 1\}$ 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 Φ^\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 Ψ^\top . Measuring the proximity between these columns in the ℓ^2 sense gives rise to the linear assignment problem

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

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

Although the problem above admits a polynomial time solution [Kuh55], typical values for n (in the order of thou-

sands) make computing this solution prohibitively expensive in practice.

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 Ψ^\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.*,

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

$$\text{s.t. } \mathbf{P}^\top \mathbf{1} = \mathbf{1}. \quad (9)$$

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 Ψ^\top to $\mathbf{C}\Phi^\top$, or vice-versa, will in general produce different results.

Balanced nearest neighbors. In order to remove the bias, we propose to incorporate additional terms in problem (8), namely minimize

$$\|\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 + \|\mathbf{C}\Phi^\top \mathbf{Q} - \Psi^\top\|_F^2 + \lambda \|\mathbf{P} - \mathbf{Q}^\top\|_F^2$$

$$\text{s.t. } \mathbf{P}, \mathbf{Q} \in \{0,1\}^{n \times n}, \mathbf{P}^\top \mathbf{1} = \mathbf{1}, \mathbf{Q}^\top \mathbf{1} = \mathbf{1}, \quad (10)$$

where the minimization is performed w.r.t. both \mathbf{P} and \mathbf{Q} . Note that we incorporated the desired property of \mathbf{P} being a permutation matrix by a soft constraint, *i.e.*, by penalizing the difference of \mathbf{P} to \mathbf{Q}^\top . Also note that the limit of $\lambda \rightarrow \infty$ leads to a convergence of \mathbf{P} to a solution of (6) meeting (7).

Instead of solving the minimization problem (10) for increasing values of λ exactly, we determine an approximate solution by alternating minimization on \mathbf{P} and \mathbf{Q} . This leads to each subproblem being a simple nearest neighbor problem and guarantees the decrease of the objective functions.

Table 1 illustrates the matching accuracy obtained by our symmetrized nearest-neighbor method in comparison to the classical nearest-neighbor (NN) and the indicator mapping (Max) approaches. As we can see, the symmetrization improves the results of the biased NN method by 2–3%.

A probabilistic model. The analysis we provided above puts in evidence two major drawbacks, namely: 1) The linear assignment approach, the nearest-neighbor search, and our bi-directional variant, all rely on the assumption that the functional map \mathbf{C} given as input aligns well the columns of Φ^\top with those of Ψ^\top in the ℓ^2 sense. 2) None of the above approaches incorporates regularity assumptions for

# basis functions	Max	NN	Balanced NN
25	6.14	30.99	33.24
50	18.12	43.51	45.65
75	26.96	52.54	55.07

Table 1: Percentage of exact matches recovered from functional maps of increasing rank. We show average results obtained by the three methods on 45 shape pairs from the FAUST dataset (intra-class, $\sim 7K$ vertices). The above comparison shows that the proposed balancing further improves the performance of the nearest neighbor technique.

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 Ψ^\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 n components, having the columns of $\mathbf{C}\Phi^\top$ as centroids in \mathbb{R}^k . For simplicity, we assume the components have uniform weight $\frac{1}{n}$, and equal covariances σ^2 . With this choice, the GMM density function is:

$$p(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n p(\mathbf{y}|\mathbf{x}_i; \sigma^2), \quad (11)$$

where we write \mathbf{y} and \mathbf{x}_i for $i = 1, \dots, n$ to denote the columns of Ψ^\top and $\mathbf{C}\Phi^\top$ respectively, and define $p(\mathbf{y}|\mathbf{x}_i; \sigma^2) = \frac{1}{(2\pi\sigma^2)^{k/2}} \exp(-\frac{\|\mathbf{y}-\mathbf{x}_i\|^2}{2\sigma^2})$.

Now let $R_\theta: \mathbb{R}^k \rightarrow \mathbb{R}^k$ 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

$$L(\theta, \sigma^2) = - \sum_{j=1}^n \log \left(\frac{1}{n} \sum_{i=1}^n p(\mathbf{y}_j | R_\theta(\mathbf{x}_i); \sigma^2) \right). \quad (12)$$

Note that the argument that minimizes (12) can be also interpreted as the argument that minimizes the Kullback-Leibler (KL) divergence between a continuous GMM distribution (represented by $\mathbf{C}\Phi^\top$) and a mixture of Dirac distributions (represented by Ψ^\top). Hence, with our probabilistic model we are choosing the distance $D(\mathbf{C}, \Psi^\top \mathbf{P})$ in Eq. (4) to be the (pseudo-)distance $D_{\text{KL}}(\mathbf{C}\Phi^\top, \Psi^\top \mathbf{P})$.

Given optimal parameters θ and σ^2 , the point-to-point

correspondence probability between \mathbf{x}_i and \mathbf{y}_j can finally be obtained as the posterior probability $p(\mathbf{x}_i|\mathbf{y}_j) = \frac{1}{n} p(\mathbf{y}_j|R_\theta(\mathbf{x}_i); \sigma^2) / p(\mathbf{y}_j)$.

Transformation function. The above probabilistic model leaves some freedom for the specific choice of a transformation R_θ . A simple example is choosing this transformation as a simple rotation, parametrized by θ . To be more flexible we instead propose to consider the general transformation:

$$R_{\mathbf{V}}(\mathbf{X}) = \mathbf{X} + \mathbf{V}, \quad (13)$$

where parameters \mathbf{V} assume the meaning of a displacement field, and $\mathbf{X} \equiv \mathbf{C}\Phi^\top$. With this definition, the overall behavior of the transformation can be controlled by regularizing \mathbf{V} .

Assuming that the given functional map \mathbf{C} represents a reasonably good matching, the refinement procedure should not be allowed to perturb the initial alignment significantly. Here we adopt the Tikhonov regularizer $\|\Gamma\mathbf{V}\|^2$ proposed in [YG89,MS10], where Γ is a low-pass operator promoting smoothly changing velocity vectors, hence coherent motion. In the inset we illustrate a smooth velocity field with coherent correspondences (blue) and a mismatch produced by simple nearest-neighbors (red). We finally obtain the regularized objective:

$$L(\mathbf{V}, \sigma^2) + \lambda \|\Gamma\mathbf{V}\|^2, \quad (14)$$

where $\lambda > 0$ is a trade-off parameter between likelihood and regularity (set to $\lambda = 3$ in our experiments).

General alignment problems like (14) 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. $\{\theta, \sigma^2\}$ and $p(\mathbf{x}_i|\mathbf{y}_j)$ 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. (14).

In Figure 2 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. As a visual measure of map quality we employ the technique described in [OBCCG13], which allows to identify the problematic areas induced by a given functional map. The quality of a map can be judged by the smoothness of the associated plots, with better maps having a more localized behavior.

The output of the EM algorithm is a set of optimal transformation parameters and a left-stochastic correspondence $\mathbf{P}^* \in \{0, 1\}^{n \times n}$. Note that similarly to the nearest-neighbor approach, we also drop the constraint $\mathbf{P}\mathbf{1} = \mathbf{1}$ in Eq. (5).

4. Point-wise map refinement

As a general representation for shape correspondences, functional maps can be adopted to compactly encode (via Eq. (3))




Figure 2: A few iterations of the minimization process for problem (4). The curve shows the average geodesic error induced by the point-wise correspondence obtained at each iteration. We visualize the quality of the initial and final maps by using the visualization technique of [OBCCG13]. In this example, the ground-truth map exhibits a similar behavior to the one we show at the last iteration.

dense point-to-point maps obtained by any matching algorithm. Therefore, one can consider improving a given functional map \mathbf{C} with the help of its associated point-to-point map, after this has been recovered with any of the methods described in the previous section. Naturally, one can repeat such a strategy and iterate between updating the point-to-point correspondence and refining the functional map. The output of such a procedure is an improved version of the initial point-to-point correspondence.

Such an iterative procedure was considered in [OBSCS*12], in which 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

$$\min_{\mathbf{C} \in \mathbb{R}^{k \times k}} \|\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 \quad (15)$$

$$\mathbf{C}^\top \mathbf{C} = \mathbf{I}_k, \quad (16)$$

which is an orthogonal Procrustes problem. Intuitively, this can be seen as a rigid alignment in \mathbb{R}^k between point sets Φ^\top and $\Psi^\top \mathbf{P}$ (see Figure 4(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 (16) imposed on the functional correspondence limit its applicability to a specific class of transformations, namely the class of volume-preserving isometries (see [OBSCS*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 constraint on \mathbf{C} and consider the problem

$$\min_{\mathbf{R} \in \mathbb{R}^{k \times k}} \|\mathbf{R}\mathbf{C}\Phi^\top - \Psi^\top \mathbf{P}\|_F^2 + \beta J^{\mathbf{C}}(\mathbf{R}), \quad (17)$$

for a regularization functional $J^{\mathbf{C}}$ which could encourage $\mathbf{R}\mathbf{C}$




Figure 3: Solving problem (17) 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.

to correspond to a smooth transformation, or could require \mathbf{R} to be a small perturbation of the identity. While for the specific example

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

problem (17) coincides with the rigid alignment problem arising from the constraint (16), a less restrictive choice for the regularization functional $J^{\mathcal{C}}$ makes the method suitable for recovering functional maps for non-isometric shape matching problems.

In our experiments we found that when (17) is combined with our proposed regularized probabilistic model for recovering the point-to-point correspondence, it is sufficient to simply update \mathbf{R} in a least-squares sense: Even without additional regularization, \mathbf{R} is determined to be a perturbation of the identity as illustrated in Figure 3. The fact that \mathbf{C} is refined in a *non-rigid* fashion can improve the refinement results significantly in comparison to orthogonal updates of \mathbf{C} as illustrated in 3d in Figure 4.

5. Experimental evaluation

We compare our iteratively refined probabilistic point-wise map recovery method with the iterative refinement procedure of Ovsjanikov *et al.* [OBBS*12] (denoted as ICP, see Section 4), which is to the best of our knowledge the only existing alternative to date. Both algorithms were implemented in Matlab/MEX and executed (single-core) on an Intel i7-3770 3.4GHz CPU with 32GB memory.

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 the two methods. We plot cumulative curves showing the percent of matches which have geodesic error smaller than a variable threshold.

We evaluate the two methods quantitatively on the FAUST [BRLB14] dataset, and qualitatively on the




Figure 4: The refinement process of a rank- k functional map can be seen as the alignment of two point sets in \mathbb{R}^k . In the first 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 (15); (c) Solution obtained with our approach (17). In the second row we show the color-coded point-wise matches obtained with each method when $k = 50$.

TOSCA [BBK08] and KIDS [RRBW*14] datasets. 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.

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 regions on the two shapes. The region correspondence is established using the ground truth, while the sets of regions are computed using the consensus technique of [RRBC14]. This way, we simulate a matching process that provides reasonably good solutions for further refinement.

We show comparisons both in the near-isometric and inter-class settings. In the former case, we use 45 pairs of humans in different poses from the intra-class subset of FAUST. All shapes have $n = 6890$ points, and the functional map is computed with $k = 30$ basis functions. Results are reported in Figure 5 (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 further 10% improvement on average.

In Figure 5 (right) we show the same curves for the non-isometric case (inter-class matching). In this case, orthogonal refinement cannot be applied due to the different properties of the input functional maps, which relate shapes under




Figure 5: Left: Comparisons with nearly-isometric shapes; in this case, orthogonal refinement (ICP) already provides accurate point-to-point recovery, which can be further improved by applying our algorithm. Right: Comparisons with non-isometric shapes. In this case, ICP cannot be applied due to the lack of orthogonality in \mathbf{C} .

non volume-preserving transformations. Additional qualitative comparisons are shown in Figure 6.

Rank. In a second set of experiments, we evaluate the capability of the different methods 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 $\mathbf{C} = \Psi^T \mathbf{P} \Phi$, 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 repeatable 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 Figure 7. As the number k of basis functions used on the two shapes (*i.e.*, rank of \mathbf{C}) increases, so does the amount of exact correspondences recovered by each method. This is also true for the simple indicator mapping approach (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).

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 Figure 2, 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 typical case where $n = 10,000$ and $k = 50$, our method takes on average 1 min. to converge, while ICP using efficient search structures (kd-trees) adds up to ~ 3 sec.

It should be noted, however, that while we employed an off-the-shelf implementation of the minimization algorithm, this code can be easily parallelized and optimized in several



Figure 6: 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 map error is visualized as a heat map encoding distance to the ground truth, growing from white to black. Note how the two methods perform comparably well in the near-isometric case (cat model), but orthogonal refinement yields large errors with more general deformations.

ways. In particular, a GPU implementation of our method remains a practical possibility.

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 demonstrated what simplifications can be used to arrive at the nearest neighbor approach and showed how simply mitigating the asymmetry present in the standard conversion procedure can lead to consistent improvements of 2-3% in accuracy. We then 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. The experimental results showed that the proposed approach yields up to 20% accuracy improvements under non-isometric deformations, reaching up to 75% *exact* point-to-point matches under good initializations.

The main limitation of our probabilistic method lies in the fact that – similar to the nearest-neighbors approach – 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




Figure 7: Percentage of exact correspondences (bold 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 (thin curves).

divergence. Extending the ideas of the symmetrized nearest-neighbor approach to the probabilistic model for removing this bias represents a possibility. Second, while for simplicity we only considered pairs of shapes related by a bijection, our method can be modified to deal with shapes having different resolutions, as well as partially similar shapes and entire shape collections (joint refinement). We believe these topics to represent exciting directions of future research.

Acknowledgments

We thankfully acknowledge Zorah Löhner and Federico Tombari for useful discussions. ER is supported by an Alexander von Humboldt fellowship. MM and DC are supported by the ERC Starting Grant “ConvexVision”.

References

- [ADK13] AFLALO Y., DUBROVINA A., KIMMEL R.: Spectral generalized multi-dimensional scaling. *CoRR abs/1311.2187* (2013). URL: <http://arxiv.org/abs/1311.2187>. 3
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 6
- [BcPP98] BURKARD R., ÇELA E., PARDALOS P., PITSOULIS L.: The quadratic assignment problem. In *Handbook of Combinatorial Optimization* (1998), Springer, pp. 241–338. 3
- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *TPAMI* 14, 2 (Feb. 1992), 239–256. 5
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (June 2014). 6
- [CR00] CHUI H., RANGARAJAN A.: A feature registration framework using mixture models. In *Proc. MMBIA* (2000), pp. 190–197. 5
- [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. 5
- [GT15] GASPARETTO A., TORSELLO A.: A statistical model of Riemannian metric variation for deformable shape analysis. In *Proc. CVPR* (2015), 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.: Robust point set registration using Gaussian mixture models. *TPAMI* 33, 8 (2011), 1633–1645. 5
- [KBB*13] KOVNATSKY A., BRONSTEIN M., BRONSTEIN A., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Comput. Graph. Forum* 32, 2pt4 (2013), 439–448. 2, 3
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proc. CVPR* (2015), pp. 905–914. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG* 30, 4 (2011), 79. 6
- [Kuh55] KUHN H. W.: The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1–2 (March 1955), 83–97. 3
- [MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *TPAMI* 32, 12 (2010), 2262–2275. 5
- [NVT*14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. *Computer Graphics Forum* 33, 5 (2014), 1–10. 3
- [OBCCG13] OVSJANIKOV M., BEN-CHEN M., CHAZAL F., GUIBAS L.: Analysis and visualization of maps between shapes. *Comput. Graph. Forum* 32, 6 (Sept. 2013), 135–145. 5
- [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, 5, 6
- [PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum* 32, 2pt4 (2013), 459–468. 2
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1 (1993), 15–36. 3
- [RCB*15] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *CoRR abs/1506.05274* (2015). URL: <http://arxiv.org/abs/1506.05274>. 2
- [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 (July 2013), 72:1–72:12. 2
- [RRBC14] RODOLÀ E., ROTA BULÒ S., CREMERS D.: Robust region detection via consensus segmentation of deformable shapes. *Computer Graphics Forum* 33, 5 (2014). 6
- [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. 2, 6
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proc. SIGGRAPH* (1995), ACM, pp. 351–358. 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. 5
- [VKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 1
- [WHG13] WANG F., HUANG Q., GUIBAS L.: Image co-segmentation via consistent functional maps. In *Proc. ICCV* (2013). 2
- [YG89] YUILLE A. L., GRZYWACZ N. M.: A mathematical analysis of the motion coherence theory. *International Journal of Computer Vision* 3, 2 (1989), 155–175. 5