

Bayesian Inference of Bijective Non-Rigid Shape Correspondence

Matthias Vestner¹, Roee Litman², Alex Bronstein^{2,3,5}, Emanuele Rodolà^{1,4} and Daniel Cremers¹

¹Technische Universität München, Germany

²Tel-Aviv Univeristy, Israel

³Technion, Israel Institute of Technology, Israel

⁴Università della Svizzera italiana, Switzerland

⁵Perceptual Computing Group, Intel, Israel

July 13, 2016

Abstract

Many algorithms for the computation of correspondences between deformable shapes rely on some variant of nearest neighbor matching in a descriptor space. Such are, for example, various point-wise correspondence recovery algorithms used as a post-processing stage in the functional correspondence framework. In this paper, we show that such frequently used techniques in practice suffer from lack of accuracy and result in poor surjectivity. We propose an alternative recovery technique guaranteeing a bijective correspondence and producing significantly higher accuracy. We derive the proposed method from a statistical framework of Bayesian inference and demonstrate its performance on several challenging deformable 3D shape matching.

1 Introduction

In geometry processing, computer graphics, and vision, estimating correspondence between 3D shapes affected by different transformations is one of the fundamental problems with a wide spectrum of applications ranging from texture mapping to animation [17]. These problems are becoming increasingly important due to the emergence of affordable 3D sensing technology. Of particular interest is the setting in which the objects are allowed to deform non-rigidly.

1.1 Related works

A traditional approach to correspondence problems is finding a *point-wise* matching between (a subset of) the points on two or more shapes. *Minimum-distortion methods* establish the matching by minimizing some structure distortion, which can include similarity of local features [30, 11, 6, 44], geodesic [27, 10, 12] or diffusion distances [14], or a combination thereof [40]. Windheuser et al. [43] used the thin shell elastic energy of triangles, while Zeng et al. [45] used higher-order structures. Typically, the computational complexity of such methods is high, and there have been several attempts to alleviate the computational complexity using hierarchical [36] or subsampling [39] methods. Several approaches formulate the correspondence problem as quadratic assignment and employ different relaxations thereof [41, 23, 33, 2, 12, 18]. Algorithms in this category typically produce guaranteed bijective correspondences between a sparse set of points, or a dense correspondence suffering from poor surjectivity.

Embedding methods try to exploit some assumption on the correspondence (e.g. approximate isometry) in order to parametrize the correspondence problem with a few degrees of freedom. Elad and Kimmel [16] used multi-dimensional scaling to embed the geodesic metric of the matched shapes into a low-dimensional Euclidean space, where alignment of the resulting “canonical forms” is then performed by sim-




Figure 1: Qualitative comparison of methods for pointwise correspondence recovery from a functional map. Current methods such as Nearest Neighbors (NN) and coherent point drift (CPD) suffer from bad accuracy and lack of surjectivity. Applying the proposed Bayesian estimation to either of them gives a guaranteed bijective matching with high accuracy. Left: We visualize the accuracy of the methods by transferring texture from the source shape \mathcal{X} to the target shape \mathcal{Y} . Neither Nearest Neighbors nor CPD produce bijective mappings. The lack of surjectivity is visualized by assigning a fixed color (green) to $y \notin im(\mathcal{X})$. Right: The geodesic error (distance between groundtruth and recovered match, relative to the shape diameter) induced by the matching is visualized on the target shape \mathcal{Y} .

ple rigid matching (ICP) [13, 8]. The works of [25, 37] used the eigenfunctions of the Laplace-Beltrami operator as embedding coordinates and performed matching in the eigenspace. Lipman et al. [24, 19, 20] used conformal embeddings into disks and spheres to parametrize correspondences between homeomorphic surfaces as Möbius transformations. By using locally injective flattenings, [4] achieve guaranteed bijective matching. However, the majority of the matching procedures performed in the embedding space often produces noisy correspondences at fine scales, and suffers from poor surjectivity.

As opposed to point-wise correspondence methods, *soft correspondence* approaches assign a point on one shape to more than one point on the other. Several methods formulated soft correspondence as a mass-transportation problem [26, 38]. Ovsjanikov *et al.* [29] introduced the *functional correspondence* framework, modeling the correspondence as a linear operator between spaces of functions on two shapes, which has an efficient representation in the Laplacian eigenbases. This approach was extended in several follow-up works [31, 22, 1]. A point-wise map is typically recovered from a low-rank approximation of the functional correspondence by a matching procedure in the

representation basis, which also suffers from poor surjectivity.

1.2 Main contributions

As the main contribution of this paper we see the formulation of the intrinsic map *denoising* problem: Given a set of point-wise correspondences between two shapes coming from any correspondence algorithm (for example, using one of the recovery algorithms outlined in Section 2), we consider them as a noisy realization of a latent bijective correspondence. We estimate this bijection using an intrinsic equivalent of the standard minimum squared error (MMSE) or minimum mean absolute error (MMAE) Bayesian estimators. To the best of our knowledge, despite their simplicity, these tools have not been previously used for deformable shape analysis.

We show that the considered family of Bayesian estimators leads to a linear assignment problem (LAP) guaranteeing bijective correspondence between the shapes. Despite the common wisdom, we demonstrate that the problem is efficiently solvable for relatively densely sampled shapes by means of the well-established auction algorithm [7] and a simple multi-

scale approach.

Finally, we present a significant amount of empirical evidence that the proposed denoising procedure consistently improves the quality of the input correspondence coming from different algorithms.

2 Pointwise map recovery

We start by briefly overviewing several recent techniques used for the computation of pointwise correspondences between non-rigid shapes. We focus on approaches relying on the functional map formalism merely because these techniques produce state-of-the-art results, emphasizing that the proposed algorithm can accept any point-wise correspondence as the input.

We model shapes as connected two-dimensional Riemannian manifolds \mathcal{X} (possibly with boundary) endowed with the standard measure da induced by the volume form. Shape \mathcal{X} is equipped with the symmetric Laplace-Beltrami operator $\Delta_{\mathcal{X}}$, generalizing the notion of Laplacian to manifolds. The manifold Laplacian yields an eigen-decomposition $\Delta_{\mathcal{X}}\phi_i = \lambda_i\phi_i$ 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 $L^2(\mathcal{X})$. Due to the latter property, any function $f \in L^2(\mathcal{X})$ can be represented via the (manifold) Fourier series expansion

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

where we use the standard manifold inner product $\langle f, g \rangle_{\mathcal{X}} = \int_{\mathcal{X}} fg da$.

Consider two manifolds \mathcal{X} and \mathcal{Y} , and let $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ be a bijective mapping between them. In [29] it was proposed to consider an operator $T : L^2(\mathcal{X}) \rightarrow L^2(\mathcal{Y})$, mapping functions on \mathcal{X} to functions on \mathcal{Y} via the composition $T(f) = f \circ \pi^{-1}$. This simple change in paradigm remarkably allows to identify maps between manifolds as linear operators (named *functional maps*) between Hilbert spaces. Because T is a linear operator, it admits a matrix representation with respect to a choice of bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ on $L^2(\mathcal{X})$ and $L^2(\mathcal{Y})$, respectively. Assuming the bases to be orthogonal, the matrix \mathbf{C} with

the elements $(\mathbf{C})_{ij} = \langle T(\phi_i), \psi_j \rangle_{\mathcal{Y}}$ provides a representation of T . In particular, by choosing the delta functions supported on the shape vertices as basis functions, one obtains a permutation $\mathbf{\Pi}$ as a matrix representation for the functional map.

A more compact way to represent T in matrix form is obtained by taking the Laplacian eigenfunctions $\{\phi_i\}_{i \geq 1}$, $\{\psi_i\}_{i \geq 1}$ of the respective manifolds as the choice for a basis. In case π is a (near) isometry, the equality $\psi_i = \pm\phi_i \circ \pi^{-1}$ holds (approximately) for all $i \geq 1$, leading to the matrix representation \mathbf{C} being diagonally dominant, i.e., $(\mathbf{C})_{ij} = \langle T(\phi_i), \psi_j \rangle_{\mathcal{Y}} \approx \pm\delta_{ij}$.

With this choice, Ovsjanikov et al. [29] proposed to truncate the matrix \mathbf{C} after the first $k \times k$ coefficients as a low-pass approximation of the functional map (typical values for k are in the range 20 – 300). This is especially convenient for correspondence problems, where one is required to solve for k^2 . At the same time, in analogy to classical Fourier analysis, the truncation has a *blurring* effect on the correspondence. As a result, recovering the original bijection π from the spectral coefficients \mathbf{C} leads to a non-trivial inverse problem.

Assume shapes \mathcal{X} and \mathcal{Y} have n points each, and let the matrices $\mathbf{\Phi}, \mathbf{\Psi} \in \mathbb{R}^{n \times k}$ contain the first $k \ll n$ eigenvectors of the respective Laplacians. For the sake of simplicity we assume $\mathbf{\Phi}$ and $\mathbf{\Psi}$ to be area-weighted, allowing us to consider the standard dot product in all equations. The expression for $\mathbf{C} \in \mathbb{R}^{k \times k}$ can now be compactly written as

$$\mathbf{C} = \mathbf{\Psi}^T \mathbf{\Pi} \mathbf{\Phi}. \quad (2)$$

Note that the matrix \mathbf{C} is now a rank- k approximation of T . The pointwise map recovery problem [32], which is highly underdetermined, consists in finding a $n \times n$ permutation $\mathbf{\Pi}$ satisfying (2). The following techniques have been proposed for this purpose.

Linear assignment problem (LAP). If we assume $k = n$, the terms on either side of (2) have the same rank and the relation can be straightforwardly inverted to yield $\mathbf{\Pi} = \mathbf{\Psi} \mathbf{C} \mathbf{\Phi}^T$. Since in the truncated setting we have $k \ll n$, the best possible solution in the l^2 sense can be obtained by looking for a permu-

tation $\mathbf{\Pi}$ minimizing $-\langle \mathbf{\Pi}, \mathbf{\Psi} \mathbf{C} \mathbf{\Phi}^T \rangle_F$. This leads to the equivalent linear assignment problem:

$$\min_{\mathbf{\Pi} \in \{0,1\}^{n \times n}} \|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{\Pi}\|_F^2 \quad (3)$$

$$\text{s.t. } \mathbf{\Pi}^T \mathbf{1} = \mathbf{1}, \mathbf{\Pi} \mathbf{1} = \mathbf{1}. \quad (4)$$

The equality between the two expressions comes from the observation that $\|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{\Pi}\|_F^2 = \|\mathbf{C} \mathbf{\Phi}^T\|_F^2 + \|\mathbf{\Psi}^T\|_F^2 - 2\langle \mathbf{\Psi} \mathbf{C} \mathbf{\Phi}^T, \mathbf{\Pi} \rangle$ for permutation matrices $\mathbf{\Pi}$. Minimizing with respect to $\mathbf{\Pi}$, and recalling that $\mathbf{\Psi}^T \mathbf{\Psi} = \mathbf{I}$, yields the equivalence.

The problem above admits an intuitive interpretation. Denoting by \mathbf{e}_i the indicator vector having the value 1 in the i th position and 0 otherwise, we see that each column of $\mathbf{\Phi}^T$ contains the spectral coefficients $\mathbf{\Phi}^T \mathbf{e}_i$ of delta functions $\delta_{x_i} : \mathcal{X} \rightarrow \{0, 1\}$ for $x_i \in \mathcal{X}$ and $i = 1, \dots, n$. Hence, the image via T of all indicator functions on \mathcal{X} is given by the columns of $\mathbf{C} \mathbf{\Phi}^T$. Problem (3) seeks for a permutation $\mathbf{\Pi}$ minimizing the distance between columns of $\mathbf{C} \mathbf{\Phi}^T$ and columns of $\mathbf{\Psi}^T \mathbf{\Pi}$ in a ℓ^2 sense.

Nearest neighbors. In [29] the authors proposed to recover a pointwise correspondence between \mathcal{X} and \mathcal{Y} by solving the nearest-neighbor problem

$$\min_{\mathbf{P} \in \{0,1\}^{n \times n}} \|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{P}\|_F^2 \quad (5)$$

$$\text{s.t. } \mathbf{P}^T \mathbf{1} = \mathbf{1}. \quad (6)$$

This can be seen as a simplified version of the LAP where the bi-stochasticity constraints (4) are relaxed, including all (binary) column-stochastic matrices \mathbf{P} in the feasible set. A global solution to (5) can be obtained in an efficient manner by solving for each column of \mathbf{P} separately: It is sufficient to seek for the nearest column of $\mathbf{C} \mathbf{\Phi}^T$ with respect to each column of $\mathbf{\Psi}^T$.

An immediate consequence of this separable approach is that its minimizers are not guaranteed to be bijections. A balanced version of (5), obtained by exchanging the roles of $\mathbf{C} \mathbf{\Phi}^T$ and $\mathbf{\Psi}^T$ in an alternating fashion was proposed in [32], with moderate increase in accuracy.

Iterative closest point (ICP). In [29] it was additionally proposed to solve for the (not necessarily bijective) \mathbf{P} according to the nearest-neighbor approach (5), followed by a refinement of \mathbf{C} via the orthogonal Procrustes problem:

$$\min_{\mathbf{C} \in \mathbb{R}^{k \times k}} \|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{P}\|_F^2 \quad (7)$$

$$\text{s.t. } \mathbf{C}^T \mathbf{C} = \mathbf{I}. \quad (8)$$

The \mathbf{P} - and \mathbf{C} - steps are alternated until convergence. In analogy to classical Iterative Closest Point (ICP) refinement [13, 8] operating in \mathbb{R}^3 , this can be seen as a rigid alignment between point sets (columns of $\mathbf{\Phi}^T$ and $\mathbf{\Psi}^T \mathbf{P}$ in \mathbb{R}^k).

Coherent point drift (CPD). The orthogonal refinement of (5), (7) assumes the underlying map to be area-preserving [29], and is therefore bound to fail in case the two shapes are non-isometric. Rodolà et al. [32] proposed to consider the non-rigid counterpart, for a given \mathbf{C} :

$$\min_{\mathbf{P} \in [0,1]^{n \times n}} D_{\text{KL}}(\mathbf{C} \mathbf{\Phi}^T, \mathbf{\Psi}^T \mathbf{P}) + \lambda \|\mathbf{\Omega}(\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{P})\|^2 \quad (9)$$

$$\text{s.t. } \mathbf{P}^T \mathbf{1} = \mathbf{1}, \quad (10)$$

where D_{KL} denotes the Kullback-Leibler divergence between probability distributions, $\mathbf{\Omega}$ is a low-pass operator promoting smooth velocity vectors, and $\lambda > 0$ controls the regularity of the assignment. Problem (9) can be seen as a Tikhonov regularization of the displacement field relating the two sets of spectral coefficients, with a measure of proximity given by the KL divergence between the two. The problem is then solved via expectation-maximization by the coherent point drift algorithm [28].

3 Bayesian map estimation

We describe a Bayesian formulation of bijective map estimation that views the given correspondence as a realization of a random process adding noise to a latent ideal correspondence. We denote by $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ the latent bijective correspondence between the

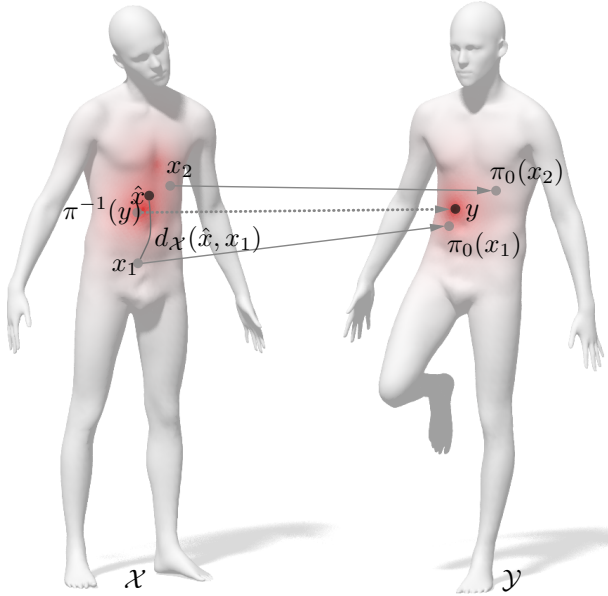


Figure 2: Conceptual illustration of the proposed Bayesian estimator. For a fixed point y , a Gaussian probability measure on \mathcal{Y} is pulled back to a measure on \mathcal{X} by the given correspondence π_0 . The estimate \hat{x} of the latent preimage $\pi^{-1}(y)$ of y is computed by minimizing the expectation of $d_{\mathcal{X}}^p(\hat{x}, \cdot)$ with respect to that measure.

shapes. Let X denote a random point on \mathcal{X} drawn from a uniform distribution, in the sense that for every measurable set $A \subset \mathcal{X}$, $\mathbb{P}(X \in A) \propto \text{area}(A)$. Given $X = x$, we denote by the conditional random variable $Y|X = x$ a point on \mathcal{Y} with a Gaussian distribution with some variance σ^2 centered at $\pi(x)$ that accounts for the uncertainty in the map. The Gaussian distribution is interpreted in the sense that for every measurable set $B \subset \mathcal{Y}$,

$$\mathbb{P}(Y \in B|X = x) \propto \int_B \exp\left(-\frac{d_{\mathcal{Y}}^2(y, \pi(x))}{2\sigma^2}\right) da(y).$$

Using Bayes' theorem, we can express the probability density of the conditional random variable $X|Y$ as

$$\begin{aligned} f_{X|Y}(x|y) &= \frac{f_{Y|X}(y|x)f_X(x)}{f_Y(y)} \propto f_{Y|X}(y|x) \\ &\propto \exp\left(-\frac{d_{\mathcal{Y}}^2(y, \pi(x))}{2\sigma^2}\right). \end{aligned}$$

Given some (possibly noisy and not necessarily bijective) correspondence $\pi_0 : \mathcal{X} \rightarrow \mathcal{Y}$, we consider $y = \pi_0(x)$ as a realization of $Y|X = x$ for every $x \in \mathcal{X}$. Our goal is to estimate the bijection π or its inverse π^{-1} from these data.

Let us fix some $y \in \mathcal{Y}$. A Bayesian estimator of $x = \pi^{-1}(y)$ given the observations π_0 can be expressed as

$$\begin{aligned} \hat{x}(y) &= \arg \min_{\hat{x}} \mathbb{E}_{X|Y=y} d_{\mathcal{X}}^p(X, \hat{x}) \\ &= \arg \min_{\hat{x}} \int_{\mathcal{X}} d_{\mathcal{X}}^p(x, \hat{x}) \exp\left(-\frac{d_{\mathcal{Y}}^2(y, \pi_0(x))}{2\sigma^2}\right) da(x). \end{aligned}$$

In the Euclidean case, the above Bayesian estimator coincides with the minimum mean absolute error (MMAE) for $p = 1$ and the minimum mean squared error (MMSE) for $p = 2$; in both cases, it has a closed-form solution as the geometric median and the centroid, respectively. The more general case discussed here can be thought of as the intrinsic counterpart of the median and the centroid.

We estimate the whole inverse map π^{-1} by minimizing

$$\hat{\pi}^{-1} = \arg \min_{\hat{\pi}^{-1}} \int_{\mathcal{X} \times \mathcal{Y}} d_{\mathcal{X}}^p(x, \hat{\pi}^{-1}(y)) e^{-\frac{d_{\mathcal{Y}}^2(y, \pi_0(x))}{2\sigma^2}} da(x) da(y). \quad (11)$$

over all bijections $\hat{\pi}^{-1} : \mathcal{Y} \rightarrow \mathcal{X}$. Note that due to the additional constraint that $\hat{\pi}^{-1}$ has to be a bijection, the estimation cannot be done for each point y independently. We also observe that iterating the process several times consistently improves the estimated map accuracy.

Finally, we note that when π is area-preserving or, more generally, scales the metric uniformly, the estimator (11) can be equivalently rewritten in terms of π as

$$\hat{\pi} = \arg \min_{\hat{\pi} : \mathcal{X} \rightarrow \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{X}} d_{\mathcal{X}}^p(x, \xi) e^{-\frac{d_{\mathcal{Y}}^2(\hat{\pi}(\xi), \pi_0(x))}{2\sigma^2}} da(x) da(\xi)$$

It is worthwhile mentioning that while being natural, the assumption of uniform prior distribution of X on \mathcal{X} (embodied in the use of the standard area measure in the above integral) can be replaced by other measures emphasizing regions where errors are less tolerable. Also, non-Gaussian noise models may be more suitable for data coming from a specific correspondence algorithm. We defer these interesting questions to future study.

Discretization. We consider the discretization of (12). We assume the shape \mathcal{X} to be discretized at n points with the corresponding discrete area elements a_i and pairwise geodesic distance matrix $\mathbf{D}_{\mathcal{X}}$. Similarly, the shape \mathcal{Y} is discretized as the same number of points, and its pairwise distance matrix is denoted by $\mathbf{D}_{\mathcal{Y}}$.

The bijective correspondence is represented by the $n \times n$ permutation matrix $\hat{\Pi}$ sought by minimizing

$$\hat{\Pi} = \arg \min_{\Pi} \text{tr}(\Pi^T \mathbf{P} \Pi), \quad (13)$$

where \mathbf{P} is an $n \times n$ matrix with the elements

$$(\mathbf{P})_{ij} = \exp\left(-\frac{(\mathbf{D}_{\mathcal{Y}})_{\pi_0(i),j}^2}{2\sigma^2}\right),$$

and $\mathbf{\Gamma}$ is an $n \times n$ matrix with $(\mathbf{\Gamma})_{ij} = (\mathbf{D}_{\mathcal{X}})_{ij}^p a_i a_j$. Note that (13) is a linear assignment problem (LAP). For directly solving the LAP with the specific structure of the score matrix given by $\mathbf{P}\mathbf{\Gamma}$, we found the auction algorithm [7] to perform the best in practice. Its average runtime complexity is $\mathcal{O}(n^2 \log n)$,

with $\mathcal{O}(n^2)$ storage complexity if a full score matrix is used. On regular hardware, this translates to several seconds for $n \sim 2.5 \times 10^3$, which quickly grows to 20 seconds for $n = 4 \times 10^3$ and almost 10 minutes for $n = 12 \times 10^3$, taking tens of gigabytes of memory. We therefore conclude that directly solving the full LAP is practical for $n \lesssim 10^4$, and in the following section propose a multi-scale scheme that can scale to much larger numbers of points.

Another computational bottleneck stems from the computation of pairwise geodesic distances. For example, using fast marching [21] the computation requires $\mathcal{O}(n^2 \log n)$ computations and $\mathcal{O}(n^2)$ storage. While the computations can be thoroughly parallelized and executed on a GPU, reducing the complexity by orders of magnitude [42], the storage of a full distance matrix is still prohibitive for $n \gtrsim 10^4$. However, since geodesic distance maps are almost everywhere smooth with constant gradient, their approximation in a truncated harmonic basis is optimal in the ℓ_2 sense [3]. Instead of storing an $n \times n$ matrix $\mathbf{D}_{\mathcal{X}}$, we store the $k \times n$ representation coefficient matrix

$$\mathbf{A}_{\mathcal{X}} = \mathbf{\Phi}^T \mathbf{D}_{\mathcal{X}},$$

where $\mathbf{\Phi}$ contains the first $k \ll n$ eigenfunctions of the Laplace-Beltrami operator on \mathcal{X} . In order to "decompress" the i -th row of $\mathbf{D}_{\mathcal{X}}$ used in the computation of the LAP score, the corresponding column of $\mathbf{A}_{\mathcal{X}}$ is multiplied from the left by $\mathbf{\Phi}$.

It is also worthwhile mentioning that while the geodesic metric is a natural candidate to compute intrinsic distances on a manifold, the proposed estimator can work with other choices. For example, diffusion distances [14] or other approximations of the geodesic distances [15] are likely to work equally well while being better amenable both for faster computation and more compact storage.

Multiscale solution. In order to reduce the computation and storage complexity associated with the direct solution of the LAP for large values of n , we adopt a multi-scale strategy. Both shapes are discretized in a hierarchical fashion using farthest point sampling, while the distances and the harmonics are calculated at the finest scale and sub-sampled.

First, a full LAP (13) is solved at a coarse scale. The produced correspondence is interpolated to the next scale and is used as the input correspondence π_0 to the LAP. While numerous interpolation techniques exist, we found that simple nearest neighbour interpolation produces satisfactory results. At the finer scale, the space of possible bijections $\hat{\pi}(i)$ is restricted to the points falling into a fixed radius r around each $\pi_0(i)$ (note that r has to be larger than the coarse sampling radius). This is equivalent to assigning infinite score to the prohibited permutations. For a sufficiently small r , this strategy results in sparse score matrices, with density significantly lower than 1%.

4 Experiments

We start by evaluating the influence of the parameters $p \in \{1, 2\}$ and σ on the quality of the Bayesian estimator (11). We initialize with noisy correspondences coming from a nearest neighbour (5) result and evaluate on two datasets with different global scales, see Figure 3. The optimal choice of σ^2 amounts to approximately 6% of the target shapes area for both choices of p but the quality is shown to be stable in a vicinity.

We test our method recover bijections from on two types of initialization, namely *functional maps of low rank* and *sparse correspondences*.

We conduct quantitative experiments on the FAUST dataset [9] (7K vertices) and on downsampled versions of the SCAPE [5] and KIDS [35] datasets (1K vertices). As quantitative quality criteria we evaluate the geodesic errors, the run times and the lack of surjectivity of the different methods. We further show that our approach can directly tackle shapes having more then 10K vertices.

4.1 Recovery from a functional map

In this set of experiments the low rank approximation is given in terms of a functional map of different ranks in the harmonic basis. Comparisons are done against nearest neighbors (NN) (5), bijective NN (3), ICP (7) and CPD (9).




Figure 3: Dependency on parameters: We evaluate the influence of the parameters p and σ on the quality of the denoised matching. We initialize with a noisy matching coming from a nearest neighbor search in descriptor space. The evaluation is done on two datasets with different global scales, namely KIDS and SCAPE.

Approximation of the groundtruth. Here we construct the low rank functional map using the known groundtruth correspondences between the shapes. This is supposed to be the ideal input for all the competing methods. As the input to our method we use the matchings found by nearest neighbors and its bijective version. We show quantitative comparisons on 71 pairs from the SCAPE dataset (near isometric, 1K vertices) and 100 pairs from the FAUST dataset (including inter-class pairs, 7K vertices). In Figures 4 and 5 we compare the accuracy, in Figure 6 the lack of surjectivity is analyzed. We only show the performance of a single application of the Bayesian estimator yet adumbrate experiments with multiple iterations in the following sections. Even after one iteration, our method outperforms the state of the art method (9) as well in accuracy as in run time (Table 1). Even on shapes having more then 10K vertices just one iteration of the Bayesian estimator gives very good results, as can be seen in Figure 7. Memory consumption and run times however limit the direct applicability of the single-scale Bayesian estimator.

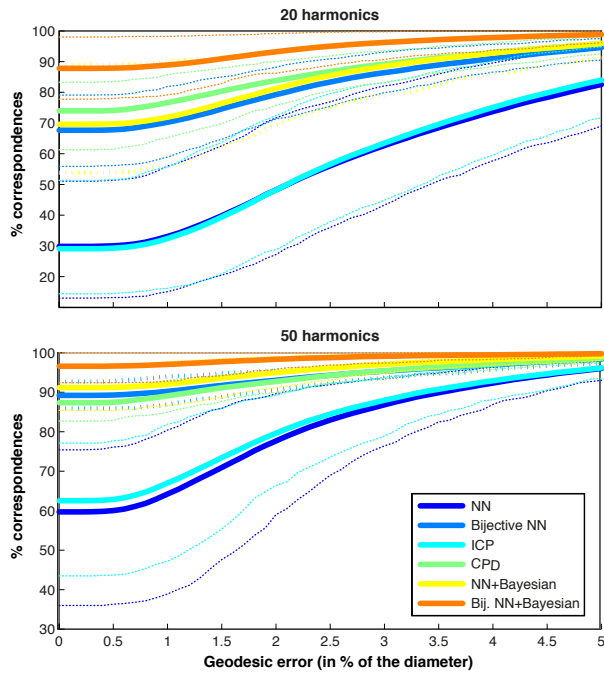


Figure 4: Recovering the groundtruth matching from a functional map with rank deficiency. We matched 70 pairs from the near-isometric SCAPE dataset (1K). Plotted are the histograms of geodesic errors (solid line: mean; 90% of the matched pairs produce results between the dotted lines). All methods boost their quality with increasing numbers of eigenfunctions. Denoising the results of nearest neighbors (yellow) outperforms the state of the art method (green) while having only a fraction of its runtime (Table 1). Even better results are achieved when initializing the Bayesian estimator with the result of bijective NN (orange).

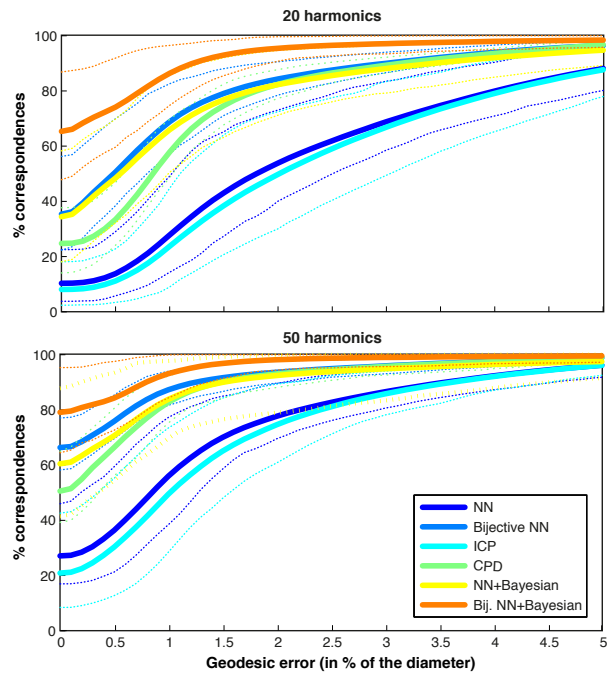


Figure 5: Evaluations from Figure 4 repeated on the higher resolution FAUST shapes. A single iteration of our denoising algorithm boosts the performance of the simple nearest neighbor approach above the state of the art (green). Notice that in this scenario bijective NN is giving better results than CPD; denoising this matching leads to 80% exact matches averaged over all pairs.




Figure 6: Histogram of distances on the target shape to the image of the source shape calculated on SCAPE (left) and FAUST (right) for matching with 20 (solid) and 50 (dotted) harmonics. In particular, the intersection with the y -axis tells the percentage of points on the target that lack a preimage. For LAP-based approaches imposing bijectivity the image of the source covers the target.

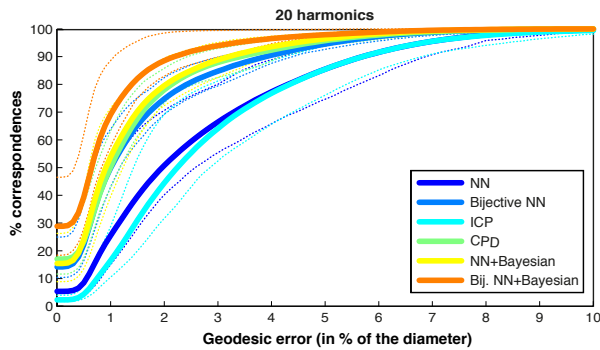


Figure 7: Our method can be applied to recover correspondences between high resolution shapes. Although the percentage of exact hits decreases with the increase of the sampling density, the geodesic errors are compelling. Due to the significant increase in runtime particularly of CPD this experiment was only done on a subset of the pairs from the SCAPE 12.5K benchmark.

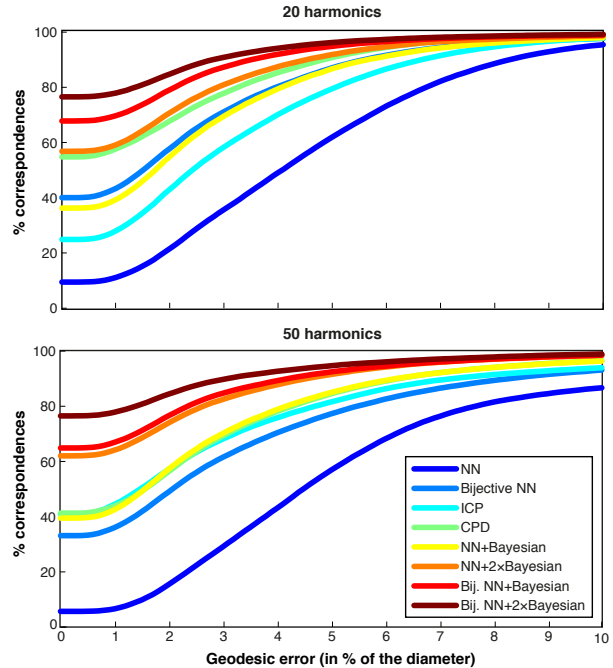


Figure 8: Evaluation on realistic input data. We test different recovery methods on a low rank functional map coming from an optimization process. Again our method gives the best results. Iterating the Bayesian estimator improves the performance even more.

Using a functional map coming from an optimization process. We follow the approach from [31] to construct a realistic functional map matching, which typically requires region features as input. These features were detected using the consensus-segmentation method proposed in [34], and the resulting regions were matched by intersection w.r.t. the ground-truth. Both of these two methods were executed with the same parameters as in their publicly available implementation. In Figure 8 this initialization is evaluated on the SCAPE dataset.

4.2 Recovery from a sparse correspondence

In this set of experiments the low rank approximation is given in terms of sparse correspondences be-

	n	1000	1000	6890	6890
	k	20	50	20	50
Nearest neighbors		0.04	0.06	1.35	2.88
Bijjective NN		2.79	2.30	463.66	253.03
ICP		0.14	0.24	12.72	30.08
CPD		4.79	4.67	1745.06	2085.65
NN + Bayesian		1.75	1.28	382.86	244.10
Bij. NN + Bayesian		4.06	3.44	746.00	440.94

Table 1: Average runtimes in seconds. We compare the runtimes of different recovery methods. Given the rank k of a functional map approximating the correspondence between shapes sampled at n points each, we report the time it takes to obtain a dense matching. See Figures 4-8 for evaluations of accuracy. Notice that while linear assignment problems are known to be time demanding to solve for larger numbers of variables, the most dramatic increase of run time occurs when applying CPD.

tween high resolution shapes. This type of input can for instance be obtained by minimizing energies under l^1 -constraints, such as [33], or appears in multi-resolution settings. We make use of groundtruth correspondences of a few points and interpolate the matching with the technique described in the caption of Figure 9. Figure 10 illustrates how iterations of the Bayesian estimator improve the matching.

5 Conclusion

We considered the problem of bijective correspondence recovery by means of denoising a given set of matches coming from any of the existing algorithms (including those not guaranteeing bijection, or producing sparse correspondences). Viewing the denoising as a Bayesian estimation problem, we formulated the intrinsic equivalent of the mean and median filters frequently employed in signal processing, with the additional constraint of bijectivity embodied through an LAP.

We find surprising the fact that such a simple idea demonstrates a consistent improvement in the correspondence quality in all experiments we have conducted. We believe that tools from estimation theory




Figure 9: Correspondence accuracy on the KIDS 1K (top plot) and FAUST 7K (middle and bottom plots). Dense matches were produced by nearest-neighbor interpolation from 20 (top and middle) and 50 (bottom) sparse matches and used as the initialization of the Bayesian estimator iterated up to five times.




Figure 10: Obtaining a dense and bijective matching of high quality by iterative application of the Bayesian estimator on a sparse correspondence. The sparse correspondence is interpolated by assigning each point the on the source shape the same image as its nearest neighbor from the sparse set (second column). This induces large geodesic errors for points being far away from the sparse set (bottom row, left) and is far from being surjective (second row from the bottom, left). Each iteration of the Bayesian estimator (five rightmost columns) yields better results, as illustrated by the magnified fragments.

that have been heavily used in other domains of science and engineering might be very useful in shape analysis, and invite the community to further explore this direction.

Of special interest are the choice of the loss function in the posterior expectation (which in this paper was restricted to the absolute and squared distance), the prior distribution of X (which we assumed uniform), and the noise distribution (which was assumed Gaussian). Alternative estimators making use of Bayesian statistics, such as maximum a posteriori (MAP) estimators, should also be explored.

References

- [1] Y. Aflalo, A. Dubrovina, and R. Kimmel. Spectral generalized multidimensional scaling. *IJCV*, 2016.
- [2] Yonathan Aflalo, Alexander Bronstein, and Ron Kimmel. On convex relaxation of graph isomorphism. *PNAS*, 112(10):2942–2947, 2015.
- [3] Yonathan Aflalo and Ron Kimmel. Spectral multidimensional scaling. *PNAS*, 110(45):18052–18057, 2013.
- [4] Noam Aigerman, Roi Poranne, and Yaron Lipman. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG)*, 33(4):69, 2014.
- [5] Dragomir Anguelov et al. SCAPE: Shape completion and animation of people. *TOG*, 24(3):408–416, 2005.
- [6] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV*, 2011.
- [7] Dimitri P Bertsekas. *Network optimization: continuous and discrete models*. Citeseer, 1998.
- [8] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *PAMI*, 14(2):239–256, 1992.
- [9] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR*, 2014.
- [10] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.
- [11] Michael M. Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010.
- [12] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *Proc. ICCV*, 2015.
- [13] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proc. Conf. Robotics and Automation*, 1991.
- [14] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.
- [15] Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. on Graphics*, 32(5):152, 2013.
- [16] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, 2001.
- [17] O.V. Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 20:1–23, 2010.
- [18] I. Kezurer, S. Kovalsky, R. Basri, and Y. Lipman. Tight relaxations of quadratic matching. *Computer Graphics Forum*, 34(5), 2015.
- [19] Vladimir G. Kim, Yaron Lipman, Xiaobai Chen, and Thomas A. Funkhouser. Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum*, 29(5):1689–1700, 2010.

- [20] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011.
- [21] Ron Kimmel and James A Sethian. Computing geodesic paths on manifolds. *PNAS*, 95(15):8431–8435, 1998.
- [22] A. Kovnatsky, M. M. Bronstein, X. Bresson, and P. Vandergheynst. Functional correspondence by matrix completion. In *Proc. CVPR*, 2015.
- [23] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proc. ICCV*, 2005.
- [24] Y. Lipman and I. Daubechies. Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics*, 227(3):1047 – 1077, 2011.
- [25] Diana Mateus, Radu P. Horaud, David Knosow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008.
- [26] Facundo Mémoli. Gromov-Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics*, pages 1–71, 2011.
- [27] Facundo Mémoli and Guillermo Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005.
- [28] A. Myronenko and X. Song. Point set registration: Coherent point drift. *TPAMI*, 32(12):2262–2275, 2010.
- [29] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. on Graphics*, 31(4):30, 2012.
- [30] Maks Ovsjanikov, Quentin Mérigot, Facundo Mémoli, and Leonidas Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.
- [31] Jonathan Pokrass, Alexander M Bronstein, Michael M Bronstein, Pablo Sprechmann, and Guillermo Sapiro. Sparse modeling of intrinsic correspondences. In *Computer Graphics Forum*, volume 32, pages 459–468. Wiley Online Library, 2013.
- [32] E. Rodolà, M. Moeller, and D. Cremers. Pointwise map recovery and refinement from functional correspondence. In *Proceedings Vision, Modeling and Visualization (VMV)*, Aachen, Germany, 2015.
- [33] Emanuele Rodolà, Alex M Bronstein, Andrea Albarelli, Filippo Bergamasco, and Andrea Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, 2012.
- [34] Emanuele Rodolà, Samuel Rota Bulò, and Daniel Cremers. Robust region detection via consensus segmentation of deformable shapes. In *Computer Graphics Forum*, volume 33, pages 97–106. Wiley Online Library, 2014.
- [35] Emanuele Rodolà, Samuel Rota Bulò, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense non-rigid shape correspondence using random forests. In *Proc. CVPR*, 2014.
- [36] Y. Sahillioğlu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum*, 30(5):1461–1470, 2011.
- [37] Alon Shtern and Ron Kimmel. Matching lbo eigenspace of non-rigid shapes via high order statistics. *arXiv:1310.4459*, 2013.
- [38] Justin Solomon, Andy Nguyen, Adrian Butscher, Mirela Ben-Chen, and Leonidas Guibas. Soft maps between surfaces. In *Computer Graphics Forum*, volume 31, pages 1617–1626, 2012.

- [39] A. Tevs et al. Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum*, 30(2):543–552, 2011.
- [40] Lorenzo Torresani, Vladimir Kolmogorov, and Carsten Rother. Feature correspondence via graph matching: Models and global optimization. In *Proc. ECCV*, 2008.
- [41] Shinji Umeyama. An eigendecomposition approach to weighted graph matching problems. *PAMI*, 10(5):695–703, 1988.
- [42] Ofir Weber, Yohai S Devir, Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Parallel algorithms for approximation of distance maps on parametric surfaces. *ACM Trans. on Graphics*, 27(4):104, 2008.
- [43] T. Windheuser, U. Schlickewei, F. R. Schmidt, and D. Cremers. Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *IEEE International Conference on Computer Vision (ICCV)*, 2011.
- [44] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009.
- [45] Y. Zeng et al. Dense non-rigid surface registration using high-order graph matching. In *Proc. CVPR*, 2010.