

A Dual Iterative Refinement Method for Non-rigid Shape Matching

Rui Xiang

Department of Mathematics,
UC, Irvine

xiangr1@uci.edu

Rongjie Lai

Department of Mathematics
Rensselaer Polytechnic Institute

lair@rpi.edu

Hongkai Zhao

Department of Mathematics
Duke University

zhao@math.duke.edu

Abstract

In this work, a robust and efficient dual iterative refinement (DIR) method is proposed for dense correspondence between two nearly isometric shapes. The key idea is to use dual information, such as spatial and spectral, or local and global features, in a complementary and effective way, and extract more accurate information from current iteration to use for the next iteration. In each DIR iteration, starting from current correspondence, a zoom-in process at each point is used to select well matched anchor pairs by a local mapping distortion criterion. These selected anchor pairs are then used to align spectral features (or other appropriate global features) whose dimension adaptively matches the capacity of the selected anchor pairs. Thanks to the effective combination of complementary information in a data-adaptive way, DIR is not only efficient but also robust to render accurate results within a few iterations. By choosing appropriate dual features, DIR has the flexibility to handle patch and partial matching as well. Our comprehensive experiments on various data sets demonstrate the superiority of DIR over other state-of-the-art methods in terms of both accuracy and efficiency.

1. Introduction

Nonrigid shape matching is one of the most basic and important tasks in computer vision and shape analysis, e.g., shape registration, comparison, recognition, and retrieval. Different from 2D images, shapes are usually represented as 2-dimensional manifolds embedded in \mathbb{R}^3 . Many challenges for nonrigid shape matching come from embedding ambiguities, i.e. shapes sharing the same metric (isometry) or very similar metrics (nearly isometry) can have drastically different xyz-coordinates representations in \mathbb{R}^3 .

To overcome these representation ambiguities, one of the commonly used strategies is to extract features which are isometrically invariant and robust to small perturbations. Along with the idea, many successful and popular approaches are based on spectral geometry [40, 31,

52, 11, 28, 44]. Theoretically, the Laplace-Beltrami (LB) operator is isometrically invariant. As a generalization of Fourier basis functions from Euclidean domains to manifolds, the eigensystem of LB operator provides complete intrinsic spectral information of the underlying manifold. Moreover, from lower eigen-modes to higher eigen-modes, LB eigenfunctions also provide a multi-scale characterization of the underlying manifold from coarse to fine resolution. Although using spectral geometry removed possible non-rigid embedding ambiguities, new ambiguities emerge in the spectral domain due to non-uniqueness of the LB eigen-system, e.g., sign ambiguity for eigenfunctions, the ambiguity of choosing a basis for the LB eigen-space corresponding to a non-simple LB eigenvalue (due to symmetry), the ambiguity of ordering for close eigenvalues (due to small perturbations). To handle these ambiguities and use spectral features accurately and robustly, a proper linear transformation (a rigid transformation for exact isometry) needs to be found to align the spectral modes between two shapes first. This linear transformation is typically computed through some matching/correlation based on given (prior) correspondence, e.g., landmarks [35, 2, 29]. Moreover, to resolve fine details and acquire accurate correspondence between two shapes, high eigen-modes need to be used. However, the use of higher eigen-modes will not only require more computation costs but, more importantly, can also cause instability with respect to small perturbations.

To tackle the instability issue of using high eigen-modes directly, one natural multi-scale approach is to start from a correspondence at a coarse scale using a few low modes and iteratively refine the correspondence at a finer scale by adding more higher modes gradually. The main motivation is that the linear transformation (a small matrix) on a coarse scale between two truncated spectral functional spaces spanned by a few low eigen-modes can be determined efficiently and stably from an initial approximate or limited correspondence. Once low eigen-modes are aligned well, an improved correspondence, especially between smooth parts of the two shapes, is likely obtained. The improved correspondence is then used to determine the

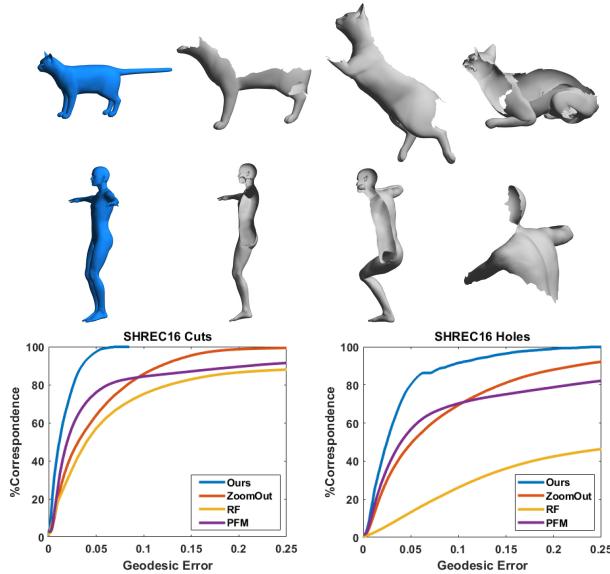


Figure 1. *Top*: First row shows examples from SHREC'16 holes, and second row shows examples from SHREC'16 cuts. We map the partial shapes in gray to full shapes in blue. *Bottom*: Geodesic error on SHREC'16 cuts and holes data set with comparison to other state-of-the-art methods.

linear transformation for the next iteration which involves more LB eigen-modes. Such a multi-scale idea for shape correspondence has been proposed in [29] for multi-scale registration using rotation-invariant sliced-Wasserstein distance and in [34] as a Zoom-out process. However, for the above straightforward multi-scale approach in the spectral domain, there are two key issues. First, in each iteration, the indiscriminate determination of the linear transformation between the truncated spectral embedding of two shapes using current correspondence of all points, many of which are incorrect, might be problematic. In Theorem 1, we show that the linear transformation between two spectral functional spaces determined using all points from an inaccurate correspondence will most likely lead to errors. These errors could be very significant and can cause either a failure for later refinement or slow convergence shown in the supplementary material. The other issue is the lack of a systematic and data-adaptive way to determine how many eigen-modes can be aligned accurately and stably by the current correspondence. Thus, it is hard to decide the appropriate jump in the number of eigen-modes for the next refinement after each iteration to achieve fast convergence. Previously, an increment of one mode was typically used in Zoom-Out [34], while a prefixed sequence of eigenmodes was proposed in [29] based on the rule of thumb.

In this work, we propose a simple and efficient dual iterative refinement strategy. This method combines complementary information, such as spatial and spectral, or local and global, to simultaneously refine the correspondence between approximately isometric shapes in an effective way.

More specifically, spatial matching via local mapping distortion (see Definition 1) is applied point-wisely to choose well-matched correspondence – anchor pairs, at the current stage. Once anchor pairs are selected, they are used to find 1) the maximal dimension of spectral functional space that can be robustly and accurately determined by these anchor pairs based on their distribution, and 2) the linear transformation that aligns the spectral basis at the next scale which will lead to a more refined correspondence at next stage. This remarkable simple strategy addresses the aforementioned two critical issues in previous multi-scale approaches using only spectral matching and allows one to utilize all well-matched pairs from the current step to jump to the next step in an accurate, efficient, stable, and data-adaptive way. We use extensive numerical experiments to show that our simple strategy, DIR, outperforms start-of-the-art model-based methods in terms of both accuracy and efficiency markedly. By choosing appropriate and application-specific dual features, DIR enjoys the flexibility to handle different scenarios, such as raw point clouds, patch and partial matching (as illustrated in Figure 1). We defer a detailed explanation of this Figure in Section 5.

Contributions. We summarize our main contributions:

- By combining spatial matching via local mapping distortion and spectral matching via the LB eigen-basis, the proposed DIR goes from coarse to fine resolution stably and effectively in a data-adaptive way which overcomes the main challenges in spectral matching.
- By generalizing this dual refinement strategy to other local and global features such as local mapping distortion and global geodesic distance, DIR handles patch and partial matching.

The rest of this paper is organized as follows. In Section 2, we provide a brief review of some related work. Then we present our proposed method, DIR, in detail followed by a discussion about a few possible extensions of our method in Section 3. After that, we conduct extensive experiments on various benchmark data sets and compare DIR to other state-of-the-art methods in terms of both accuracy and efficiency in Section 5. We conclude the paper in Section 6.

2. Related work

Designing effective shape descriptors is crucial for shape registration or correspondence. In general, descriptors can be categorized as pointwise or pairwise. Extrinsic pointwise descriptors [49, 21, 15, 41] are easy to compute but usually not very accurate, especially when a non-rigid transformation is involved. To handle non-rigid transformation, different intrinsic descriptors either in spatial domain, such as geodesics distance signatures [53], heat kernel signatures [47] and wave kernel signatures [4], or in

spectral domain using the eigensystem of the LB operator [31, 52, 11, 28, 35, 29], are proposed. Then various nearest neighbor searching or linear assignment methods are used in the descriptor space to find the dense point correspondence. For point-wise descriptors in the spectral domain, such as functional map [35], a proper linear transformation is needed to align the spectral basis to remove ambiguities in the eigensystem of LB operator. The determination of the linear transformation itself is typically based on some prior knowledge, e.g., given landmarks, region correspondences, or orientation preserving properties [39], and can be challenging to achieve good accuracy especially when the dimension of the spectral embedding is high, which is needed to resolve fine shape features. On the other hand, spatial domain pointwise descriptors are usually defined smoothly on shapes and that makes it difficult to provide accurate correspondences. Our method aims at tackling these issues raised in aligning shapes using pointwise descriptors.

Intrinsic pairwise descriptors, such as pairwise geodesic distance [7, 10, 55] and kernel functions [32, 46, 54] have also been proposed for non-rigid shape correspondence problems. Although the matching of pairwise descriptors is stricter and may be more robust and accurate, it requires to solve a quadratic assignment problem (QAP) which is NP-hard. Various methods have been proposed to solve the QAP approximately in a more computational tractable way e.g sub-sampling [48], coarse-to-fine [56], geodesic distance sparsity enforcement methods [19] and various relaxation approaches [1, 8, 25, 30, 13, 16]. One popular approach is to relax the nonconvex permutation matrix (representing pointwise correspondence) constraint in the QAP to a doubly stochastic matrix (convex) constraint [1, 13]. However, both the pairwise descriptors and the doubly stochastic matrix are dense matrices, which make the relaxed QAP still challenging to solve even for a modest size problem. Recently, a novel local pairwise descriptor and a simple, effective iterative method to solve the resulting quadratic assignment through sparsity control was proposed in [57] for shape correspondence between two approximately isometric surfaces.

3. Dual Iterative Refinement Method

3.1. Functional map

Spectral geometry is widely used in shape analysis [40, 31, 52, 11, 28, 35, 29, 45, 22]. Leveraging the spectral representation, functional map formalism is introduced in [35] to solve non-rigid shape correspondence problem. It improves earlier point-based spectral methods [23, 24, 33, 36] which directly match the spectral embeddings of shapes. We provide a brief introduction to the concept of functional maps which is used in our multi-scale DIR process.

Given (\mathcal{M}_1, g_1) as a closed 2-dimensional Riemannian manifold, the LB operator uniquely determines the underlying manifold up to isometry [6]. Eigen-functions of LB operator form an orthonormal basis on the underlying manifold and can be used as intrinsic and multi-scale descriptors for shapes. In discrete setting, the manifold \mathcal{M}_1 is usually represented as a triangulated mesh with vertices $\{\mathbf{x}_i\}_{i=1}^n$. The LB matrix is given by $\mathcal{L}_{\mathcal{M}_1} = \mathcal{A}_1^{-1} \mathcal{W}_1$ [37], where \mathcal{A}_1 is the diagonal element area matrix of \mathcal{M}_1 and \mathcal{W}_1 is the standard cotangent weight matrix. The discrete truncated k -dimensional spectral embedding of \mathcal{M}_1 can be expressed as a matrix $\Phi_{\mathcal{M}_1}^k \in \mathbb{R}^{n \times k}$ whose rows are the first k LB eigen-functions evaluated at each point. Similarly, we can define the first k LB eigen-embedding for \mathcal{M}_2 as $\Phi_{\mathcal{M}_2}^k$.

The key idea of functional map is to lift a shape correspondence $T : \mathcal{M}_1 \rightarrow \mathcal{M}_2$ to a linear functional map, between the function spaces $\mathcal{C}(\mathcal{M}_1)$ and $\mathcal{C}(\mathcal{M}_2)$. The functional map for a given correspondence T can be represented as a linear transformation between two given bases of $\mathcal{C}(\mathcal{M}_1)$ and $\mathcal{C}(\mathcal{M}_2)$, which becomes a matrix when the shapes are discretized and the bases are truncated. In practice, LB eigen-functions are commonly used as the basis for the function space. For example, given a permutation matrix Π representing a point-to-point map T from \mathcal{M}_1 to \mathcal{M}_2 , the functional map C is defined as:

$$C = \arg \min_C \|\Phi_{\mathcal{M}_1} C - \Pi \Phi_{\mathcal{M}_2}\|_F^2 \quad (1)$$

Theoretically, if \mathcal{M}_1 and \mathcal{M}_2 are isometric, then the corresponding LB eigen-functions are the same up to possible ambiguities caused by sign switch and non-simple eigenvalues, which belong to an orthonormal group. This motivates the following constrained functional map [29]

$$\min_{C \in \mathfrak{R}} \|\Phi_{\mathcal{M}_1} C - \Pi \Phi_{\mathcal{M}_2}\|_F^2 = \text{Proj}_{\mathfrak{R}}(\Phi_{\mathcal{M}_1}^\top \Pi \Phi_{\mathcal{M}_2}) \quad (2)$$

where $\mathfrak{R} = \{C \mid C^\top C = Id\}$ denotes the set of orthonormal matrices and the projection to \mathfrak{R} is provided as $\text{Proj}_{\mathfrak{R}}(A) = UV^\top$ with the singular value decomposition (SVD) of $A = U\Sigma V^\top$. Since this paper aims at tackling nearly isometric shape matching problem, we adopt (2) which is equivalent to (1) for the isometric case.

Typically, the functional map is obtained via a least square optimization with various constraints and regularizations, such as preservation of given landmarks, commutativity with LB operator, and sparsity [46, 39]. Once the optimized functional map C is computed, for any $p \in \mathcal{M}_1$, the correspondence $T(p)$ can be obtained by solving the following spectral matching problem:

$$T(p) = \arg \min_{q \in \mathcal{M}_2} \|\Phi_{\mathcal{M}_1}(p)C - \Phi_{\mathcal{M}_2}(q)\|_F^2 \quad (3)$$

However, the determination of an accurate functional map C can be difficult when limited prior knowledge or a

poor correspondence is used. On the one hand, it becomes harder when more eigen-modes are involved since the degrees of freedom of the functional map grow quadratically in terms of the eigen-modes involved and, moreover, high eigen-modes are less computationally stable with respect to perturbations or noises. On the other hand, confining to low eigen-modes limits resolution of the spectral representation and hence the accuracy of shape correspondence. A natural multi-scale idea is to start the functional map from a coarse resolution involving a few low eigen-modes and construct a coarse correspondence T . Then the coarse correspondence is used to help determining the functional map at a finer scale and hence a more refined correspondence. This process can be iterated until a fine enough resolution is achieved. The main motivation is that the functional map at a coarse scale – a small matrix, can be determined efficiently and stably from an initial approximate or limited correspondence, which is then used to improve the correspondence at the next iteration. This strategy leads to iteratively computing (2) and (3), or (1) and (3), which is the key idea proposed in [29] using rotation-invariant sliced Wasserstein distance and in the Zoom-out process [34]. As a crucial component of these iterative refinement methods, the new functional map C is obtained from the current correspondence of all points. Since intermediate correspondence, especially at the beginning, can be quite inaccurate, this naive strategy may lead to significant errors.

Next, we provide both theoretical evidences on the effect of correspondence error of C in (2). This can be problematic for any iterative refinement procedure based on spectral geometry. Assume we have two perfectly isometric manifolds $\mathcal{M}_1, \mathcal{M}_2$ and their corresponding discrete truncated spectral embeddings $\Phi_{\mathcal{M}_1}, \Phi_{\mathcal{M}_2}$. Without loss of generality, we assume Id is the ground truth correspondence between \mathcal{M}_1 and \mathcal{M}_2 (otherwise, we can shuffle row vectors of $\Phi_{\mathcal{M}_2}$ according to the ground truth correspondence). The ground truth functional map is an orthonormal matrix $C_T \in \mathfrak{R}$, i.e. $\Phi_{\mathcal{M}_2} = \Phi_{\mathcal{M}_1} C_T$ and $C_T = \text{Proj}_{\mathfrak{R}}(\Phi_{\mathcal{M}_1}^\top \Phi_{\mathcal{M}_2})$.

Theorem 1 Given $\Phi_{\mathcal{M}_2} = \Phi_{\mathcal{M}_1} C_T$ with $C_T \in \mathfrak{R}$, let us assume a one-to-one correspondence Π is an inaccurate correspondence which maps a portion of $\Phi_{\mathcal{M}_1}$ accurately to the corresponding part of $\Phi_{\mathcal{M}_2}$, while the remaining part of Π is inaccurate. Without loss of generality, we write $\Phi_{\mathcal{M}_1} = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$ and $\Phi_{\mathcal{M}_2} = \begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix}$ where $X_1, Y_1 \in \mathbb{R}^{n_1 \times k}, X_2, Y_2 \in \mathbb{R}^{n_2 \times k}$ and $n_1 + n_2 = n$. We let $\Pi \Phi_{\mathcal{M}_2} = \begin{pmatrix} Y_1 \\ \sigma Y_2 \end{pmatrix}$ for a permutation matrix $\sigma \in \mathbb{R}^{n_2 \times n_2}$.

Let $C_a = \arg \min_{C \in \mathfrak{R}} \|\Phi_{\mathcal{M}_1} C - \Pi \Phi_{\mathcal{M}_2}\|_F^2$. Then the spectral norm $\|C_a - C_T\|_2 > 0$ with probability at least $1 - \eta$ with $\eta = \sum_{j=0}^{\lfloor n_2/2 \rfloor} \frac{1}{2^j j!(n_2 - 2j)!}$.

Notice that η decreases rapidly as n_2 grows. For example, when $n_2 = 25$, $\eta \approx 10^{-12}$. We provide a detailed proof of this theorem in the supplementary material where we also numerically verify that using points from an inaccurate correspondence may lead to significant errors in the functional map. This can cause failure for iterative refinement.

Another critical issue for multi-scale approach in the spectral domain is how to iteratively increase the resolution, i.e., the dimension of spectral embedding, in an accurate, stable and data-adaptive way. Most approaches [54, 29, 34] just adopt an empirical or prefixed increasing sequence.

Motivated from the above difficulties and limitations of iterative refinement approaches in the spectral domain, we propose a dual iterative refinement method that iteratively updates on the spatial and spectral domains. The first key idea is to choose well-matched pairs using local spatial matching from current correspondence, called anchor pairs, and only use them to determine the functional map which helps to construct a much more improved correspondence for the next iteration. In order to choose high-quality anchor pairs from a given correspondence, we zoom in at each corresponding pair and measure local mapping distortion. This step integrates local spatial information in the spectral refinement process. The second key idea is to find the maximal dimension of spectral embedding that can be robustly and accurately determined by these anchor pairs according to their distribution in the spectral domain using singular value analysis on their correlation matrix.

3.2. Local mapping distortion

In order to choose well-matched anchor pairs, or to filter out bad correspondences, to compute the functional map C more accurately, we use the following *local mapping distortion* (LMD) introduced in [57] to measure the matching quality of a given correspondence pair and choose those pairs with small LMD. Without relying on any information of the ground truth correspondence, this distortion provides a quantitative measurement to check the accuracy of the map through its continuity and local distance preservation.

Definition 1 (Local mapping distortion) Let $T : \mathcal{M}_1 \rightarrow \mathcal{M}_2$ be a map between two manifolds. For any point $\mathbf{x} \in \mathcal{M}_1$, consider its γ -geodesic ball in \mathcal{M}_1 as $\mathcal{B}_\gamma(\mathbf{x}) = \{\mathbf{y} \in \mathcal{M}_1 \mid d_{\mathcal{M}_1}(\mathbf{x}, \mathbf{y}) \leq \gamma\}$. LMD of T at \mathbf{x} is defined as:

$$\mathcal{D}_\gamma(T)(\mathbf{x}) = \frac{1}{|\mathcal{B}_\gamma(\mathbf{x})|} \int_{\mathbf{y} \in \mathcal{B}_\gamma(\mathbf{x})} DE_T(\mathbf{x}, \mathbf{y}) d\mathbf{y} \quad (4)$$

where $DE_T(\mathbf{x}, \mathbf{y}) = |d_{\mathcal{M}_1}(\mathbf{x}, \mathbf{y}) - d_{\mathcal{M}_2}(T(\mathbf{x}), T(\mathbf{y}))|/\gamma$ and $|\mathcal{B}_\gamma|$ is the volume of \mathcal{B}_γ .

Based on the definition of LMD, it is straightforward to check if T is an isometric map, then $\mathcal{D}_\gamma(T)(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{M}_1, \gamma > 0$. Conversely, if $\mathcal{D}_\gamma(T)(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{M}_1$ for

some $\gamma > 0$, then T is isometric. Numerically, equation (4) can be approximated by:

$$\mathcal{D}_\gamma(T)(\mathbf{x}_i) \approx \frac{\sum_{\mathbf{x}_j \in \mathcal{B}_\gamma(\mathbf{x}_i)} \mathcal{A}_1(j) DE_T(\mathbf{x}_i, \mathbf{x}_j)}{\left(\sum_{\mathbf{x}_j \in \mathcal{B}_\gamma(\mathbf{x}_i)} \mathcal{A}_1(j) \right)} \quad (5)$$

where \mathcal{A}_1 is the area element of the mesh representing \mathcal{M}_1 . By normalizing the distance with respect to γ , LMD is ro-

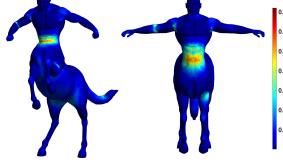


Figure 2. LMD on ground truth correspondence.

bust to global scaling as well as local sampling variation. Figure 2 shows that, for the ground truth correspondence, LMD is very small almost everywhere except in regions where non-isometric distortion is large. Hence the proposed LMD also serves as a good unsupervised error metric when ground truth is not available.

3.3. Our Method

We first precompute K leading LB eigenfunctions for each shape. Here K can be determined by computation cost limitation, the noise level in the data, or desired accuracy of the correspondence. In any case, it should not exceed what the mesh resolution can support for the discretized shape, i.e., a few mesh points are needed in each nodal domain.

Our method can start with an initial correspondence provided by SHOT [49] based on extrinsic point-wise features, or a few given landmarks used as initial anchor pairs which can be fixed or updated in later refinement, or by any other (fast but not necessarily very accurate) methods. Then we start DIR which includes the following three simple steps: ① Choose anchor pairs from current correspondence using LMD criterion; ② Determine the spectral dimension and the corresponding functional map based on the selected anchor pairs; ③ Update the correspondence using the current functional map. In addition, we enforce two stopping criteria, the total number of iterations and the spectral dimension supported by anchor pairs reaches K , whichever is satisfied first. Here are more detailed descriptions of each step.

Step 1: Selecting anchor pairs. By setting a proper LMD threshold ϵ , the set of pairs $\{(x_\ell, T(x_\ell)) | \mathcal{D}_\gamma(T)(x_\ell) < \epsilon\}$ are selected as anchor pairs from the current correspondence which will be used to guide the next refinements. It is important to note that both the threshold ϵ and anchor pairs are updated in later refinement. By decreasing ϵ with iterations, the quality of anchor pairs is also improved.

Step 2: There are two components described as follows.

(1) **Determining the proper spectral dimension based on anchor pairs.** Given a set of anchor pairs, an important

question is to find a proper spectral dimension that can be determined accurately and stably according to the distribution of the anchor pairs in the spectral embedding. Due to possible degeneracy of the distribution, the dimension can be quite smaller than the number of anchor pairs. We use singular value decomposition (SVD) [50], a.k.a principal component analysis, to find the dimension well expanded by a set of anchor pairs in spectral domain.

$$U \Sigma V^\top = \Phi_{\mathcal{M}_1}^K (\{x_\ell\}_{\ell=1}^m)^\top \Phi_{\mathcal{M}_2}^K (\{T(x_\ell)\}_{\ell=1}^m). \quad (6)$$

We threshold the singular values to determine the proper dimension. One can use a simple thresholding strategy. In our implementation, we adopt a normalization strategy to make the threshold adaptive to the data and noise level. After normalizing all singular values by the mean of 10 largest singular values, the dimension cut is set at where the sum of ten consecutive normalized singular values is smaller than a threshold (0.1 for all our tests).

(2) **Computing functional map based on anchor pairs.** Once the proper spectral dimension, k , is determined by the selected anchor pairs $\{(x_\ell, T(x_\ell)) | \ell = 1, \dots, m\}$, we compute the functional map for k dimensional spectral embedding only based on the anchor pairs as follows:

$$\begin{aligned} C &= \arg \min_{C \in \Re} \sum_{\ell=1}^m \| \Phi_{\mathcal{M}_1}^k (x_\ell) C - \Phi_{\mathcal{M}_2}^k (T(x_\ell)) \|_F^2 \\ &= \text{Proj}_{\Re} \left((\Phi_{\mathcal{M}_1}^k (\{x_\ell\}_\ell))^\top \Phi_{\mathcal{M}_2}^k (\{T(x_\ell)\}_\ell) \right) \end{aligned} \quad (7)$$

This restricted version to compute a functional map minimizes possible corruption from inaccurate correspondence and leads to a more accurate estimation of the functional map as discussed in the Section 3.1 and supplementary material. Meanwhile, it also reduces the computation cost.

Step 3: Construct the new correspondence. Using the functional map computed from selected anchor pairs in the properly enlarged spectral embedding space, a refined correspondence is constructed by solving the assignment problem (3), where a KNN search method is applied.

We summarize the full procedure in Algorithm 1. **Computation complexity** The complexity for computing K leading eigen-vectors of a $n \times n$ sparse matrix corresponding to the discretized LB operator is $O(Kn \log n)$. The complexity for checking LMD is $O(n)$ since geodesic distance is only computed in a fixed neighborhood, e.g., first ring or second ring, at each point. To compute the functional map, the complexity of matrix multiplication is at most $O(nK^2)$ and SVD decomposition is at most $O(K^3)$ since the number of anchor pairs is at most n and the spectral dimension k is at most K which is prefixed and far less than n . KNN search used to solve Equation (3) has a complexity of $O(n \log(n))$ [51]. Hence, our method is of complexity $O(n \log(n))$ altogether. It is still $O(n \log(n))$ using

Algorithm 1: Dual Iterative Refinement (DIR)

input An initial correspondence, $\Phi_{\mathcal{M}_1}^K$ and $\Phi_{\mathcal{M}_2}^K$ (e.g. K leading LB eigenfunctions), LMD error $\{\epsilon_i\}_{i=1}^N$ thresholds and the maximum iteration N .
while $i \leq N$, or $k_i \leq K$ **do**

1. Find anchor pairs $\{(x_\ell, T(x_\ell)) \mid \mathcal{D}_\gamma(T)(x_\ell) < \epsilon_i\}_{\ell=1}^{m_i}$ via LMD criterion from current correspondence T .
2. Find spectral dimension k_i from SVD of $\Phi_{\mathcal{M}_1}^K(\{x_\ell\}_\ell)^\top \Phi_{\mathcal{M}_2}^K(T(\{x_\ell\}_\ell))$.
3. Update the functional map $C = \text{Proj}_{\mathfrak{R}} \left(\left(\Phi_{\mathcal{M}_1}^{k_i}(\{x_\ell\}_\ell) \right)^\top \Phi_{\mathcal{M}_2}^{k_i}(\{T(x_\ell)\}_\ell) \right)$
4. Update the correspondence (for all $p \in \mathcal{M}_1$) $T(p) = \arg \min_q \|\Phi_{\mathcal{M}_1}^{k_i}(p, :)C - \Phi_{\mathcal{M}_2}^{k_i}(q, :)C\|^2$.

end

other global features such as geodesic distance to anchor points since we limit the maximal number of anchor points (as the number of LB eigen-functions). The complexity of computing geodesic distance and solving the resulting assignment problem is still $O(n \log(n))$.

4. Discussion

In this section, we discuss other possible extensions of our approach and a few specific applications.

Combination of local and global features. For shape correspondence problems, an efficient and robust approach should use both local and global features. So far we have mainly talked about using spatial and spectral features, which are perfectly complementary in the sense of local and global information. Anchor pair selection by the LMD criterion is based on local spatial features while the spectral features are global. We would like to point out that the proposed strategy of DIR process can be extended to other combinations of appropriate local and global features in different applications. For examples, one may alternatively use Heat Kernel Signature [12], Wave Kernel Signature [4] or Geodesic Distance Signature (GDS) [53] as global features. For shapes with holes and boundaries, we choose GDS which is less sensitive to local mesh distortions and boundaries for most interior points. However, if accurate spectral information is available, taking advantage of the multi-scale representation in spectral embedding leads to better accuracy and efficiency according to our tests.

DIR with limited initial landmarks. In our previous discussion, the first collection of anchor pairs is selected from a given initial correspondence such as the one obtained from comparing SHOT features. Our method does enjoy the flexibility to incorporate given landmarks. This is applauded in applications like shape matching in morphological stud-

ies in medical imaging where landmarks could be annotated based on specific tasks [20]. Once a few human-annotated and required landmarks are given, which are taken as the initial anchor pairs and fixed in later iterations, DIR will converge to a stable solution. The fewer landmarks are provided, the more iterations are usually needed. The final convergence performance of our numerical tests, as shown in Figure 6, indicates that DIR can provide accurate correspondence based on only four landmarks and it is stable with respect to different initialization.

Matching of point clouds without mesh. In real applications, point clouds with well-constructed global triangulation are usually hard to obtain. For point clouds without global mesh, we use the local mesh method to compute the LB eigenvalues and LMD [27, 57]. DIR works well on matching raw point clouds as shown in section 5.

Patch/partial matching. Patch/partial matching often comes with difficulties due to artificial boundaries, different sizes and topological perturbations. Spectral information is either unavailable or unreliable to use. However, LMD is not affected by the above difficulties at interior points. By replacing spectral features with geodesic distance to anchor points, which is global and less sensitive to the above difficulties for most interior points, DIR can handle patch or partial matching well as shown in our numerical experiments.

5. Experiments

In this section, we conduct comprehensive experiments to evaluate the performance of our method on various benchmark data sets. In all experiments, no pre-processing, such as obtaining a low resolution model or pre-computing the geodesic distance matrix, is required in our algorithm. Raw point clouds (with or without meshes) are directly used. In our comparisons, all geodesic errors from existing methods are obtained from the associated error curve data appeared in the papers. Computation using ZoomOut is produced from the code on GitHub shared by the authors [34]. All experiments using our methods are conducted in Matlab with 16GB RAM and Intel i7-6800k CPU.¹

Error Metric. Suppose the constructed correspondence maps $\mathbf{x} \in \mathcal{M}_1$ to $\mathbf{y} \in \mathcal{M}_2$ while the true correspondence maps \mathbf{x} to \mathbf{y}^* . We measure the quality of our results by computing the geodesic error defined as $e(x) = \frac{d_{\mathcal{M}_2}(y, y^*)}{\text{diam}(\mathcal{M}_2)}$, where $\text{diam}(\mathcal{M}_2)$ is the geodesic diameter of \mathcal{M}_2 .

Hyperparameters. We use the second ring neighborhood for LMD criterion; maximum iteration number is 10 with the LMD threshold as $[0.26, 0.22, 0.18, 0.14, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1]$. The initial correspondence is given by the KNN search result based on SHOT features [49]. These hyperparameters,

¹Code is available at https://github.com/ruixiang440/Dual_Iterative_Refinement_Method

especially the LMD threshold, are not sensitive to different data sets. We use the same hyperparameters on TOSCA and SCAPE data sets. For SHREC’16 and patch matching, a maximum of 800 anchor pairs serve as reference points for computing geodesic features.

TOSCA [9], **SCAPE** [3]. TOSCA data set consists of 76 shapes in 8 different categories (human and animal shapes) with vertex numbers ranging from 4k to 50k. We present the results of our method using mesh structures with the maximal spectral dimension as 1000 and 500 respectively. We also apply our method using only point clouds (no mesh) with the maximal spectral dimension as 1000 (using local mesh method [27]). We perform the same experiments on SCAPE data set which has 72 shapes (12,500 vertices for each) of the same person with different poses.

In Figure 3, we compare our method with the following methods: Blended [26], Best Conformal [26], Smooth Shells [18], Kernel Marching [54], SEQA [57], RSWD [29], ZoomOut [34], Divergence-Free Shape Interpolation [17], SRFM [38] and BCICP [39]. Our method with 1000 spectral basis outperforms all methods, our method with 500 spectral basis outperforms almost all methods, and our method on purely point cloud inputs still achieves good performance. Moreover, the run time of our method DIR is also significantly less than other methods (see Table 1).

Table 1 indicates the computation efficiency of our methods by showing the average run time on several examples from TOSCA including shapes with vertices ranging from 4k to 50k. The reported time of our method counts in all steps including the computation of LB eigenfunctions. As discussed in Section 3.3, the complexity of our method is $O(n \log(n))$. But in practice, when computing shapes with more than 20,000 vertices, our computer suffers a computation speed slow-down from the vast RAM usage because of our limited RAM capacity. Hence, for Cat and David shape from TOSCA, the run time is higher than expected. Most state-of-the-art approaches have a computation complexity of $O(n^2)$ and do not report run time for shapes over 10,000 vertices. Our method is very efficient compared with state-of-the-art methods which report computation time.

Model	Wolf	Centaur	Horse	Cat	David
Number of Vertices	4344	15768	19248	27894	52565
SEQA (s)	59	531	801	929	1681
Kernel Matching (s)	60	NA	NA	NA	NA
ZoomOut(3-150 spectral basis) (s)	25	277	449	1267	4312
Consistent Shape Matching (s)	483	NA	2118	3299	NA
Our Method with 1000 spectral basis (s)	36	144	191	474	1774
Our Method with 500 spectral basis (s)	19	106	131	330	1301

Table 1. Average run time for shapes from TOSCA comparing with SEQA [57], ZoomOut [34], Consistent Shape Matching [5] (without including precomputation time for the geodesic distance matrix) and Kernel Matching [54].

SHREC’16 [14] SHREC’16 Partial Correspondence

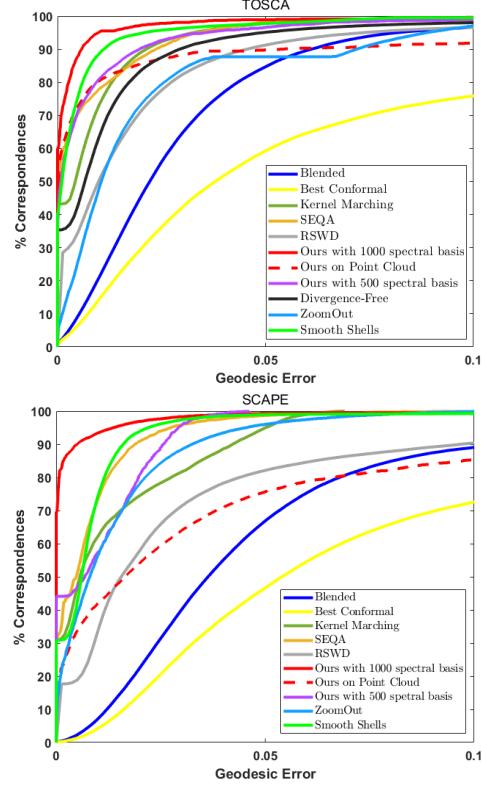


Figure 3. Geodesic error on the TOSCA and SCAPE data sets with comparison to other methods.

benchmark data set consists of 8 types of isometric human or animal shapes in different poses with regular ‘cuts’ and irregular ‘holes’. We test our method by matching each partial shape to the corresponding full shape. Since spectral basis is sensitive to mesh ‘cuts’ and ‘holes’, we use geodesic distance to anchor points as global features. We still use SHOT for the initialization. In Figure 1, we compare our method with ZoomOut [34], Partial Functional Maps [42] and Random Forests [43]. The results show that our method is quite flexible and robust to handle shape matching in different scenarios. It significantly outperforms other state-of-the-art methods and achieves the highest accuracy on this data set.

Patch Matching. Since there is no standard data set for patch matching, we report several experiments using patches cut from TOSCA data set. The first case contains two examples with boundaries, different sizes (partial patching) and topological changes simultaneously. In the first example, we take a portion of an arm (finger tips also removed) from one of two nearly isometric centaur shapes, and then map the partial arm onto the other entire centaur. In the second example, we map a body patch onto the whole shape. The original centaur is a closed surface with no holes, while the arm and the body patch are not closed and have holes and boundaries. Our method performs very well

even for this challenging example as shown in Figure 4.

The second case is matching two patches containing both overlap and non-overlap parts. We match an arm without the hand to a portion of an arm with the hand, where the forearm is the common part. Since there is no correspondence between the non-overlap parts, a post LMD test is added to prune out those points. The result is shown in Figure 5. Again our method performs quite well. These experiments indicate that our method is robust to size differences, artificial boundaries and topological changes.

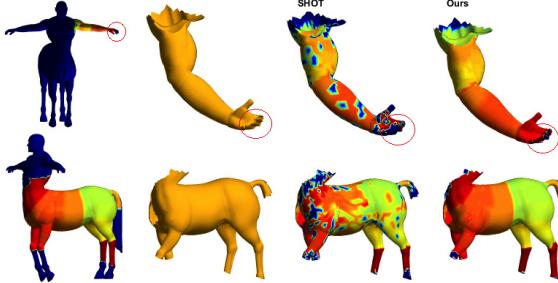


Figure 4. Partial matching between body parts in the second column and the entire body. The first column: the entire centaurs with the ground truth color map (non-blue area). Extra points are colored in blue. Removed finger tips are highlighted by red circles. The third and fourth columns: the color map of a result using SHOT and our method, respectively.



Figure 5. Patch matching. From left to right: the source patch (with the ground truth color map in the non-blue area and extra points colored in blue), the target patch, mapping result using SHOT and mapping result using our method. Blue colored regions in columns three and four indicate points not passing the post LMD test.

Experiments Using Limited Landmarks. We select 4 pairs of centaur shapes from TOSCA, initialize our method with different numbers of randomly chosen landmark pairs (without using SHOT for initialization) and then plot the average geodesic error curve. These landmark pairs are fixed and always belong to our selected anchor pairs in later iterations. After different number of iterations, the final performance is illustrated in Figure 6. Even with four initial random landmarks, the performance is excellent although more iterations are needed. It shows stability and robustness of our method with respect to initial correspondences.

Growth of Anchor Pairs and Spectral Dimensions. One key feature of our method is the selection of well matched pairs used to determine spectral dimension and the corresponding functional map at next iteration in an automatic

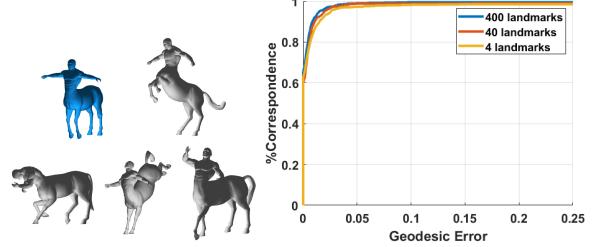


Figure 6. Average performance of our method on several pairs of centaur shapes from TOSCA (blue centaur map to gray centaurs), given different numbers of initial random annotated landmarks.

and data-adaptive way. A steady growth and improvement (by decreasing the mapping distortion threshold with iterations) of anchor pairs (updated in each iteration) means a steady growth of the spectral dimension and increase of resolution which leads to an effective and fast refinement of the correspondence. The left image in Figure 7 shows the average spectral dimensions for several classes of shapes from TOSCA or SPACE data set. The right image in Figure 7 shows the growth pattern of anchor pairs. These indicate that our method can extract anchor pairs and increase spectral resolution in a steady and efficient way.

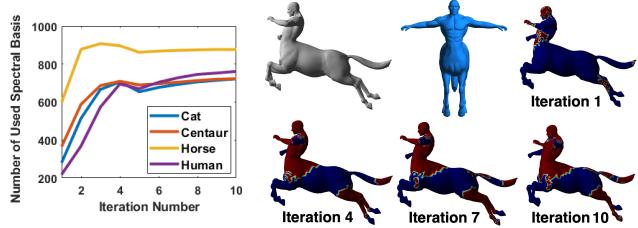


Figure 7. Left plot: The average spectral dimension (with a maximal number of LB eigenfunctions set to 1000) for cat, horse, centaur class from TOSCA and human from SCAPE at each iteration. Right figure: Selected anchor pairs (color-coded in red) during iterations of finding the correspondence between the gray centaur and the blue one. The anchor pairs detected at iteration 1 are directly from matching SHOT descriptors.

6. Conclusion

We propose a simple and efficient iterative refinement strategy utilizing dual features, spatial and spectral or local and global. By only relying on selected well-matched pairs to guide the next refinement, our proposed method combines complementary information in an optimal and data-adaptive way. Our method shows superior performance on extensive tests compared to other state-of-the-arts methods in terms of accuracy, efficiency and stability.

Acknowledgement

R. Lai's research is supported in part by an NSF Career Award DMS-1752934. H. Zhao's research is supported in part by DMS-2048877 and DMS-2012860.

References

[1] Yonathan Aflalo, Alexander Bronstein, and Ron Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015.

[2] Yonathan Aflalo, Anastasia Dubrovina, and Ron Kimmel. Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision*, 118(3):380–392, 2016.

[3] Dragomir Anguelov, Praveen Srinivasan, Hoi-Cheung Pang, Daphne Koller, Sebastian Thrun, and James Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in neural information processing systems*, pages 33–40, 2005.

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

[5] Omri Azencot, Anastasia Dubrovina, and Leonidas Guibas. Consistent shape matching via coupled optimization. In *Computer Graphics Forum*, volume 38, pages 13–25. Wiley Online Library, 2019.

[6] Pierre Bérard, Gérard Besson, and Sylvain Gallot. Embedding riemannian manifolds by their heat kernel. *Geometric & Functional Analysis GAFA*, 4(4):373–398, 1994.

[7] Alex Bronstein, Michael Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006.

[8] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006.

[9] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.

[10] Alexander M Bronstein, Michael M Bronstein, Ron Kimmel, Mona Mahmoudi, and Guillermo Sapiro. A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision*, 89(2-3):266–286, 2010.

[11] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1704–1711, 2010.

[12] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1704–1711. IEEE, 2010.

[13] Qifeng Chen and Vladlen Koltun. Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047, 2015.

[14] Luca Cosmo, Emanuele Rodolà, Jonathan Masci, Andrea Torsello, and Michael Bronstein. Matching deformable objects in clutter. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 1–10. IEEE, 2016.

[15] Anastasia Dubrovina and Ron Kimmel. Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. *Advances in Adaptive Data Analysis*, 3(01n02):203–228, 2011.

[16] Nadav Dym, Haggai Maron, and Yaron Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv preprint arXiv:1705.06148*, 2017.

[17] Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Divergence-free shape interpolation and correspondence. *arXiv preprint arXiv:1806.10417*, 2018.

[18] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12265–12274, 2020.

[19] Andrea Gasparetto, Luca Cosmo, Emanuele Rodola, Michael Bronstein, and Andrea Torsello. Spatial maps: From low rank spectral to sparse spatial functional representations. In *2017 International Conference on 3D Vision (3DV)*, pages 477–485. IEEE, 2017.

[20] Xianfeng Gu, Yalin Wang, Tony F Chan, Paul M Thompson, and Shing-Tung Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE transactions on medical imaging*, 23(8):949–958, 2004.

[21] Stefan Gumhold, Xinlong Wang, and Rob S MacLeod. Feature extraction from point clouds. In *IMR*. Citeseer, 2001.

[22] Li Han, Shuning Liu, Bing Yu, Shengsi Xu, and Rui Xiang. Orientation-preserving spectral correspondence for 3d shape analysis. *Journal of Imaging Science and Technology*, 2019.

[23] Varun Jain and Hao Zhang. Robust 3d shape correspondence in the spectral domain. In *Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on*, pages 19–19. IEEE, 2006.

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

[25] Itay Kezurer, Shahar Z Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxation of quadratic matching. In *Computer Graphics Forum*, volume 34, pages 115–128. Wiley Online Library, 2015.

[26] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, page 79. ACM, 2011.

[27] Rongjie Lai, Jiang Liang, and Hongkai Zhao. A local mesh method for solving pdes on point clouds. *Inverse Problems & Imaging*, 7(3), 2013.

[28] Rongjie Lai, Yonggang Shi, Kevin Scheibel, Scott Fears, Roger Woods, Arthur W Toga, and Tony F Chan. Metric-induced optimal embedding for intrinsic 3d shape analysis. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2871–2878. IEEE, 2010.

[29] Rongjie Lai and Hongkai Zhao. Multiscale nonrigid point cloud registration using rotation-invariant sliced-wasserstein distance via laplace–beltrami eigenmap. *SIAM Journal on Imaging Sciences*, 10(2):449–483, 2017.

[30] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1482–1489. IEEE, 2005.

[31] B. Levy. Laplace-beltrami eigenfunctions: Towards an algorithm that understands geometry. *IEEE International Conference on Shape Modeling and Applications, invited talk*, 2006.

[32] Xiuwen Liu, Arturo Donate, Matthew Jemison, and Washington Mio. Kernel functions for robust 3d surface registration with spectral embeddings. In *2008 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.

[33] Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[34] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):155, 2019.

[35] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30:1–30:11, 2012.

[36] Maks Ovsjanikov, Quentin Merigot, Facundo Memoli, and Leonidas Guibas. One point isometric matching with the heat kernel. *CGF*, 29(5):1555–1564, 2010.

[37] Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics*, 2(1):15–36, 1993.

[38] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured regularization of functional map computations. In *Computer Graphics Forum*, volume 38, pages 39–53. Wiley Online Library, 2019.

[39] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)*, 37(6), 2018.

[40] M. Reuter, F.E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design*, 38:342–366, 2006.

[41] Emanuele Rodolà, Alex M Bronstein, Andrea Albarelli, Filippo Bergamasco, and Andrea Torsello. A game-theoretic approach to deformable shape matching. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 182–189. IEEE, 2012.

[42] Emanuele Rodolà, Luca Cosmo, Michael Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. *Computer Graphics Forum*, 36(1):222–236, 2017.

[43] Emanuele Rodolà, Samuel Rota Bulò, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense non-rigid shape correspondence using random forests. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4177–4184. IEEE, 2014.

[44] Raif M Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233. Eurographics Association, 2007.

[45] Stefan C Schonsheck, Michael M Bronstein, and Rongjie Lai. Nonisometric surface registration via conformal laplace-beltrami basis pursuit. *arXiv preprint arXiv:1809.07399*, 2018.

[46] Alon Shtern and Ron Kimmel. Iterative closest spectral kernel maps. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 499–505. IEEE, 2014.

[47] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer graphics forum*, 28(5):1383–1392, 2009.

[48] Art Tevs, Alexander Berner, Michael Wand, Ivo Ihrke, and H-P Seidel. Intrinsic shape matching by planned landmark sampling. In *Computer Graphics Forum*, volume 30, pages 543–552. Wiley Online Library, 2011.

[49] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *European conference on computer vision*, pages 356–369. Springer, 2010.

[50] Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.

[51] Pravin M Vaidya. An $o(n \log n)$ algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4(2):101–115, 1989.

[52] B. Vallet and B. Levy. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum (Proceedings Eurographics)*, 2008.

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

[54] Matthias Vestner, Zorah Lähner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodolà, Alex Bronstein, Michael Bronstein, and Ron Kimmel. Efficient deformable shape correspondence via kernel matching. In *3D Vision (3DV), 2017 International Conference on*, pages 517–526. IEEE, 2017.

[55] Matthias Vestner, Roee Litman, Emanuele Rodolà, Alex Bronstein, and Daniel Cremers. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proc. CVPR*, pages 6681–6690, 2017.

[56] Chaohui Wang, Michael M Bronstein, Alexander M Bronstein, and Nikos Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 580–591. Springer, 2011.

[57] Rui Xiang, Rongjie Lai, and Hongkai Zhao. Efficient and robust shape correspondence via sparsity-enforced quadratic assignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9513–9522, 2020.