Title
Efficient and Robust Shape Correspondence Methods

Permalink
https://escholarship.org/uc/item/2g22f6jd

Author
Xiang, Rui

Publication Date
2020

UNIVERSITY OF CALIFORNIA,
IRVINE

Efficient and Robust Shape Correspondence Methods

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Mathematics

by

Rui Xiang

Dissertation Committee:
Professor Hongkai Zhao, Chair
Professor Jack Xin
Professor Long Chen

2020

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

ACKNOWLEDGMENTS

I would like to thank my advisor, Professor Hongkai Zhao, for his patient and profound advice for both research and life starting from 2014 during my exchange time at UC Irvine. He helped me transform from a young and naive university student to an independent and mature grown-up.

I would also like to thank Professor Rongjie Lai, from Rensselaer Polytechnic Institute. He and Professor Hongkai Zhao, together, guided me through many difficulties in my research. His great insight helped me avoid lots of detours.

I would like to extend my appreciation to all of my defense committee members, Professor Jack Xin and Professor Long Chen as well as advancement committee members, Professor Babak Shahbab and Professor John Lowengrub for their time, insight and expertise.

Besides, I would like to thank my family and my friends. I could never finish my Ph.D study without their great support. And my special thanks to Miss Little Six.

I would like to thank my funding supports through Professor Hongkai Zhao's NSF grants DMS-1418422, DMS-1622490 and DMS-1821010. Part of my research was performed while I was visiting the Institute for Pure and Applied Mathematics (IPAM), which was supported by the NSF DMS-1440415.

VITA

Rui Xiang

EDUCATION

Doctor of Philosophy in Mathematics **2015–2020**
University of California, Irvine *Irvine, CA*

Bachelor of Science in Since **2011–2015**
Chinese University of Hong Kong *Hong Kong, China*

RESEARCH EXPERIENCE

Graduate Research Assistant **2015–2020**
University of California, Irvine *Irvine, CA*

TEACHING EXPERIENCE

Teaching Assistant **2015–2020**
University of California, Irvine *Irvine, CA*

REFEREED JOURNAL PUBLICATIONS

Orientation-Preserving Spectral Correspondence for 3D Shape Analysis **2020**
Journal of Imaging Science and Technology

REFEREED CONFERENCE PUBLICATIONS

Efficient and Robust Shape Correspondence via Sparsity-Enforced Quadratic Assignment **Jun 2020**
Computer Vision and Pattern Recognition

ABSTRACT OF THE DISSERTATION

Efficient and Robust Shape Correspondence Methods

By

Rui Xiang

Doctor of Philosophy in Mathematics

University of California, Irvine, 2020

Professor Hongkai Zhao, Chair

In this work, we introduce two methods to handle nearly isometric shape correspondence problem. The first method,*Sparsity-Enforced Quadratic Assignment* (SEQA), introduces a novel local pairwise descriptor and then develops a simple, effective iterative method to solve the resulting quadratic assignment through sparsity control. The pairwise descriptor is based on the stiffness and mass matrix of finite element approximation of the Laplace-Beltrami differential operator. The key idea of our iterative algorithm is to select pairs with good correspondence as anchor pairs based on a new criterion we introduced called local mapping distortion, and then solve a regularized quadratic assignment problem the neighborhoods of corresponding selected anchor pairs through sparsity control. Various pointwise global features with reference to these anchor pairs can be used to improve the dense shape correspondence further. Inspired by the powerful ability of local mapping distortion on selecting accurate anchor pairs, we further introduce the second method, *Dual Iterative Refinement* (DIR). It is a simple and efficient algorithm which combines dual features, spatial and spectral, or local and global, in a complementary and optimal way. DIR first uses local spatial feature, local mapping distortion to obtain anchor pairs which are used to determine an appropriate dimension of the spectral space and the corresponding functional map. Then the functional map is used to update a new correspondence. Both methods allow us to deal with open surfaces, partial matching, and topological perturbations robustly. We use various experiments to show the efficiency, quality, and versatility of our methods on large data sets, patches, and point clouds (without global meshes).

Chapter 1

Introduction

1.1 Shape Correspondence

Geometric modeling and shape analysis is ubiquitous in computer vision [53, 49], computer graphics [79, 44], medical imaging [69], virtual reality [39, 23], 3D prototyping and printing [57], data analysi [71], etc [35, 8, 48, 68]. Shape correspondence [75, 81, 9] is a basic task in shape registration, comparison, recognition, and retrieval. Unlike images, shapes do not have a canonical representation domain or basis and do not form a linear space. Moreover, their embedding can be highly ambiguous even for intrinsically identical ones. Further complications in practice include noise, topological perturbations (holes), partial shapes, and lack of a good triangulation. These difficulties pose both modeling and computational challenges for shape modeling and analysis [15, 54].

Given two manifolds $\mathcal{M}_1$ and $\mathcal{M}_2$ sampled by two point clouds $\mathcal{P}_1 = \{\mathbf{x}_i\}_{i=1}^n$ and $\mathcal{P}_2 = \{\mathbf{y}_i\}_{i=1}^n$ respectively, the typical task of dense shape correspondence is to find a point-to-point map φ between $\mathcal{P}_1$ and $\mathcal{P}_2$ as illustrated in Figure 1.1. φ admits an injective and subjective correspondence as a permutation matrix $P \in \Pi_n = \{\Pi \in \{0,1\}^{n\times n} |\ \Pi\vec{1} = \vec{1}, \Pi^T\vec{1} = \vec{1}\}$.

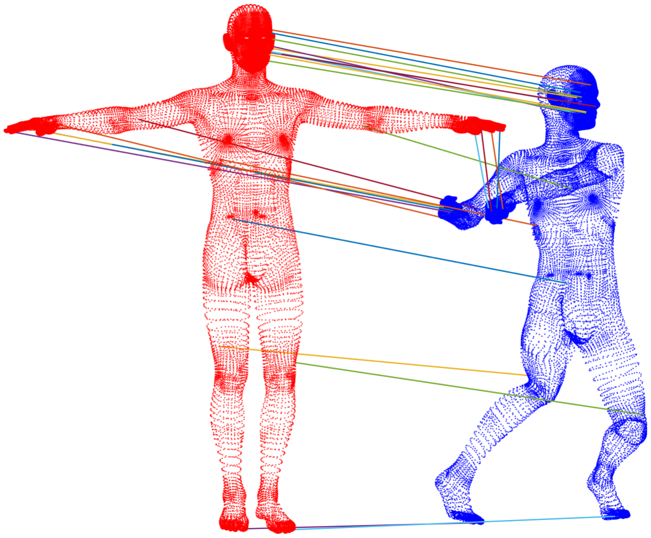

Figure 1.1: Shape correspondence of the same person of two poses.

In many applications, $\mathcal{M}_1$ and $\mathcal{M}_2$ are nearly isometric. In order to obtain an accurate correspondence, the first step of most literature is to design desirable descriptors which usually can be categorized as pointwise, or pairwise.

1.1.1 Pointwise Descriptor

Pointwise descriptor gives each point a feature vector. Suppose $\mathcal{P}_1$ and $\mathcal{P}_2$ have pointwise feature matrix $\mathcal{F}_1, \mathcal{F}_2 \in \mathbb{R}^{n \times k}$ where k is the feature dimension, $\mathbb{P}$ is usually obtained by solving the following linear assignment problem (LAP).

$$\underset{\mathbb{P}\in\Pi_n}{\arg\min} \|\mathbb{P}\mathcal{F}_1 - \mathcal{F}_2\|^2 = \underset{\mathbb{P}\in\Pi_n}{\arg\min} < \mathbb{P}, \mathcal{F}_2\mathcal{F}_1^T > \tag{1.1}$$

The induced linear assignment problem 1.1 is usually solved by auction algorithm [8] with complexity $O(n^2 log(n))$ in the descriptor space to find the dense point correspondence. Nearest neighbor searching [33] can also provide an approximated solution with complexity $O(nlog(n))$.

Most Pointwise descriptors can be categorized as extrinsic (invariant under rigid transformation) [71, 28, 20, 60], or intrinsic (invariant under isometric transformation). Extrinsic pointwise descriptors usually have difficulties producing accurate dense correspondence, especially if there is non-rigid transformation involved. Many intrinsic pointwise descriptors in the space domain, such as geodesics distance signatures [75], heat kernel signatures [67], wave kernel signatures [5], and in spectral domain using eigen-functions of the Laplace-Beltrami operator (LBO) have been proposed [59, 42, 74, 13, 38, 39]. For example, functional maps [50] aims to find proper linear combinations of truncated basis functions, e.g., first four eigen-functions of the Laplace-Beltrami operator as illustrated in Figure 1.2, based on some prior knowledge, e.g., given landmarks and/or region correspondence, as the pointwise descriptor. As a generalization of Fourier basis functions from Euclidean domains to manifolds, the eigen-system of LB operator provides a system of intrinsic features of the underlying manifold in the spectral domain. Moreover, from lower modes to higher modes, LB eigen-functions 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 eigen-functions, the ambiguity of choosing a basis for the LB eigen-space corresponding to a non-simple LB eigen-value (due to symmetry), the ambiguity of ordering for close eigen-values (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 basis between two shapes first. This linear transformation is typically computed through some matching/correlation based on given (prior) correspondence, e.g., landmarks [50, 3, 39]. High spectral modes need to be used to resolve fine details and acquire accurate correspondence between two shapes. However, the higher the spectral mode, the less robust or computationally stable it is to small perturbations.

To tackle the instability issue of using higher spectral 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 and finer scale by adding more and more, higher and higher spectral modes gradually. The primary motivation is that the linear transformation (a small matrix) on a coarse scale between two spectral spaces spanned by a few low spectral modes can be determined efficiently and stably from an initial approximate or limited correspondence. Once low mode spectral features 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 linear transformation for the next iteration, which involves more and higher spectral modes. Such a multi-scale idea for shape correspondence has been proposed in [39] for multi-scale registration using rotation-invariant sliced-Wasserstein distance and in [46] as a Zoom-out process. However, for the above straightforward multi-scale approach solely in the spectral domain, there are two key issues. First, in each iteration, the determination of the linear transformation between the spectra spaces of two shapes using current correspondence of all points, many of which are incorrect, indiscriminately might be problematic. The linear transformation between two spectral 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 appendix. The other issue is the lack of a systematic and data-adaptive way to determine how many spectral modes can be supported, given the current correspondence. Thus, it is hard to decide what is the appropriate jump in the number of spectral modes for refinement after each iteration that can achieve fast convergence. Previously, an increment of one mode was typically used to maintain convergence in Zoom-Out [46], while a prefixed sequence of eigen-modes was proposed in [39] based on the rule of thumb.

These intrinsic pointwise descriptors are typically non-local and require to solve certain partial differential equations, e.g., the Laplace-Beltrami equation, on a well-triangulated mesh. Hence they can be sensitive to topological perturbations and boundary conditions.

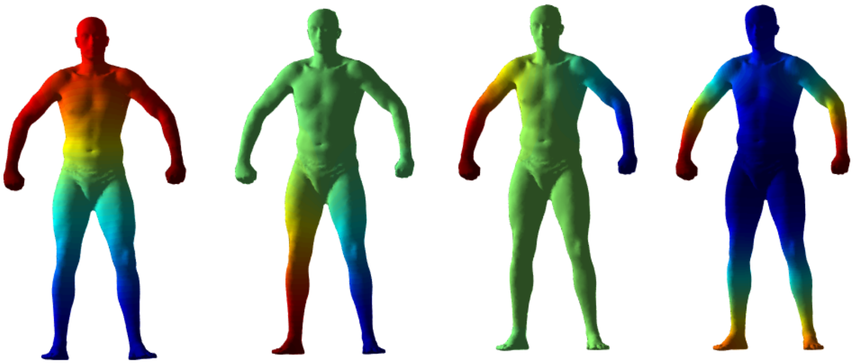

Figure 1.2: First four eigen-functions of Laplace-Beltrami operator on a shape

Moreover, pointwise descriptors based on a truncated basis in the spectral domain lose fine details in the geometry; besides, there does not exist a clear and general guidance on where to set the dimensional truncation.

1.1.2 Pairwise Descriptor

Pairwise descriptor prescribes some relation/similarity between each pair of points in the data set. Suppose $\mathcal{P}_1$ and $\mathcal{P}_2$ have pairwise feature matrix $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbb{R}^{n\times n}$, $\mathbb{P}$ is usually obtained by solving the following quadratic assignment problem (QAP) or graph matching (GM).

$$\arg\min_{\mathbb{P}\in\Pi_n} ||\mathbb{P}\mathcal{Q}_1 - \mathcal{Q}_2\mathbb{P}||^2 = \arg\min_{\mathbb{P}\in\Pi_n} < \mathbb{P}, \mathcal{Q}_2\mathbb{P}\mathcal{Q}_1 > \tag{1.2}$$

The induced quadratic assignment problem 1.2 is NP-hard [40]. There does not exist efficient and effective algorithm to solve the original NP-hard problem. Most existing methods are based on certain relaxation.

Using good pairwise descriptors, such as pairwise geodesic distance matrix [77] or kernel functions [76], to find shape correspondence is usually more robust and accurate since the matching needs to satisfy more and stricter constraints to minimize some distortion. Different kinds of methods have been proposed to solve the QAP approximately in a more computational tractable way. A sub-sampling method is introduced in [70]. Similarly, a coarse-to-fine scheme [78] divides the whole process into several steps to reduce the computation complexity. [63] explores the ability of genetic optimization to find the permutation matrix that matching two normalized geodesic distance matrix. [24] proposes a geodesic distance sparsity enforcement and combines spectral information. Other various approaches [2, 11, 34, 41, 17, 22] also aims to find an efficient and effective relaxation. One popular approach is to relax the nonconvex permutation matrix (representing pointwise correspondence) constraint in the QAP to a doubly stochastic matrix (convex) constraint [2, 17]. Based on this common relaxation, [2] studies the theoretical guarantee of relaxing permutation matrix to doubly stochastic matrix under certain conditions of the pairwise descriptors and also discusses the phenomenon of dropping non-negativity constrain. More theoretical studies are introduced in general graph matching problem [25, 18].

Other than the NP-hard nature of QAP, another challenging problem in 1.2 is that most pairwise descriptors and the permutation matrix are all dense. It is very common that a point cloud contains over 10,000 points; hence in a general setting, e.g., using geodesic distance matrix as a pairwise descriptor, one will need to precompute and store two huge hundred thousands by hundred thousands geodesic distance matrices, and a hundred thousands by hundred thousands permutation matrix. It is a tremendous challenge not only on the computation cost but also on the storage capacity.

1.2 Our Result

We briefly introduce our two methods in this section and highlight the main contributions of them. The first method is named as *Sparsity-Enforced Quadratic Assignment* (SEQA), and the second method is named as *Dual Iterative Refinement* (DIR).

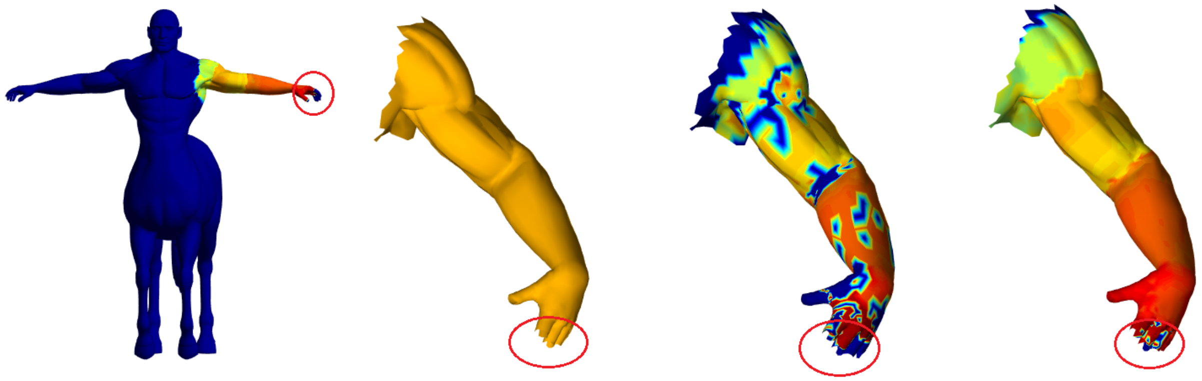

Figure 1.3: Example of partial matching with topological changes using SEQA. Topological changes are highlighted by red circles. Each patch in the second column is mapped onto the entire shape in the first column, and the non-blue area is the ground truth map. Extra points in the entire shape are colored in blue. The third column is the mapping result using SHOT, and the last column is the mapping result from our method.

1.2.1 Sparsity-Enforced Quadratic Assignment

In this work, we propose a novel approach for dense shape correspondence for two nearly isometric surfaces based on a new local pairwise descriptor and an efficient iterative algorithm with sparsity control for the doubly stochastic matrix to solve the corresponding relaxed QAP. The main novelty and contribution of our proposed method include:

- A local pairwise descriptor using the combination of the stiffness (corresponding to the finite element approximation of the LBO) and the mass matrix (corresponding to local area scaling). It only involves interactions among local neighbors, which results in the sparse nature of theses two matrices, and they are extremely simple to compute. Note that all local interactions are coupled like heat diffusion through the whole shape. In other words, global and full spectral information of LBO is embedded implicitly in our pairwise descriptor. Due to the locality, the descriptor enjoys stability and good performance for open surfaces and with respect to topological perturbations, as shown in Figure 1.3. The sparsity of the pairwise descriptor also reduces the computation and storage cost for the relaxed QAP.

- An efficient iterative algorithm with sparsity control for the resulting relaxed QAP. We first use a local mapping distortion measurement defined as local mapping distortion to select point pairs from both shapes with good correspondence as anchor pairs for the next iteration. Using regularity of the map, we enforce that the neighborhoods of an anchor pair can only map to each other which induces a sparsity structure in the doubly stochastic matrix. It results in a significant reduction of variables and hence, the computation and storage cost in each iteration. As we demonstrate in the numerical experiments, the number of high-quality anchor pairs grow quickly with iterations, which forms a virtuous cycle to improve the quality of the correspondence map gradually.

- Local mapping distortion criterion provides a quantitative measurement of the quality of each pair correspondence. Various methods, e.g., heat kernel signature, geodesic distance signature, can easily take advantage of these anchor pairs to further improve the whole mapping quality as a post-processing approach. These methods initially require human annotated landmarks, while in our algorithm, no human annotated landmark is needed.

- All of the novel ingredients in our algorithm, e.g., sparse pairwise descriptor, sparsity-enforced QAP, and local mapping distortion criterion, can be easily coupled with other algorithms to improve their performance and reduce computation cost. The combination of our local mapping distortion criterion and the well-known functional map results in our second method *Dual Iterative Refinement*.

1.2.2 Dual Iterative Refinement

In this work, we propose a simple and efficient dual iterative refinement strategy. This method combines dual information, such as spatial and spectral (or local and global), in a complementary and optimal way for dense correspondence between two nearly isometric shapes. The main novelties and contributions of our proposed method include:

- We propose to use dual information (spatial and spectral, or local and global) to select anchor pairs and to guide the iterative refinement process. The key idea is to use a zoom-in operation to measure the mapping quality at each point quantitatively based on local mapping distortion. This local spatial information (global in the spectral domain) enables us to choose anchor pairs from both shapes that are well-matched at the current stage.
- We design a data-adaptive way to determine the multi-scale path for the refinement process. Once anchor pairs are selected, they are used to 1) find the maximal dimension of spectral modes that can be robustly and accurately determined by these anchor pairs based on their distribution in spectral space (e.g., using singular value analysis on their correlation matrix), and 2) find the linear transformation that aligns spectral features at finer scale based on these anchor pairs. This remarkable simple strategy addresses the aforementioned two critical issues in previous multi-scale approaches. It allows one to optimally utilize all and only well-matched pairs in the current step to jump to the next finer scale 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.
- We explore extensions to different applications. By choosing appropriate and application-specific dual features, DIR has the flexibility to deal with different scenarios, such as raw point clouds, patch, and partial matching (as illustrated in Figure 1.4). We defer a detailed explanation of this Figure in Section 3.5.

The rest of this paper is organized as follows. In Chapter 2, we discuss SEQA in detail and illustrate corresponding experimental results; in Chapter 3, we discuss DIR in detail and illustrate corresponding experimental results.

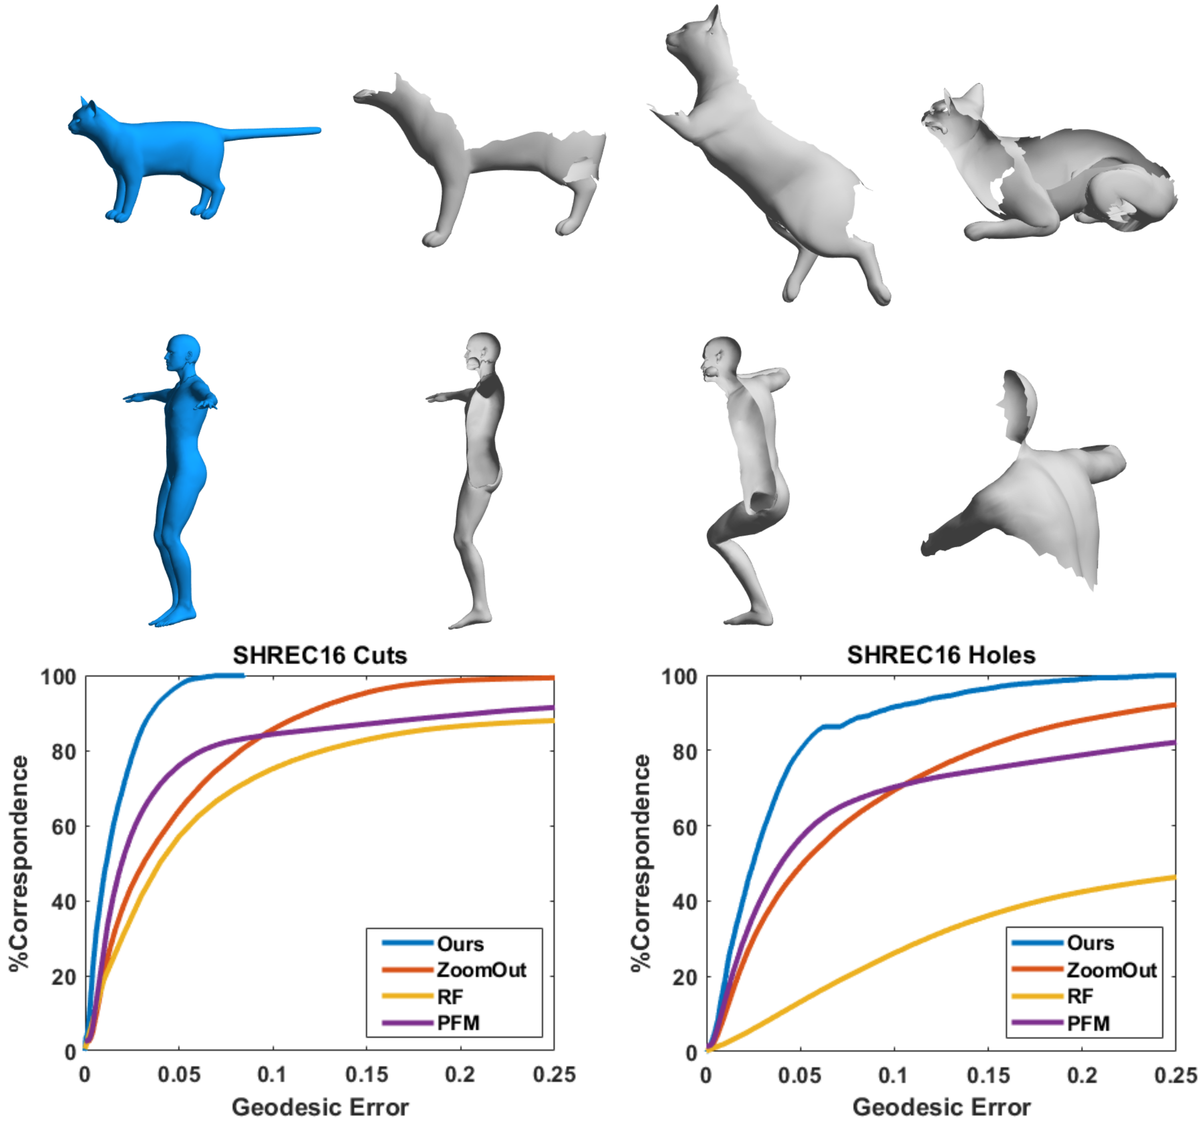

Figure 1.4: *Top:* First row shows examples from SHREC'16 holes, and second row shows examples from SHREC'16 cuts. We map the partial shapes in gray to full shapes in blue. *Bottom:* Geodesic error on SHREC'16 cuts and holes data set with comparison to other state-of-the-art methods.

Chapter 2

Sparsity-Enforced Quadratic Assignment

In this chapter, we explain our *Sparsity-Enforced Quadratic Assignment* (SEQA) method in detail. We introduce our quadratic assignment model based on a local pairwise descriptor in Section 2.1 and then present an efficient iterative algorithm to solve the quadratic assignment problem with sparsity control in Section 2.3. In Section 2.6, we extend our method to point cloud data and patch matching. Numerical experiments are demonstrated in Section 2.7 and conclusion follows.

2.1 Quadratic Assignment Model Using Local Pairwise Descriptors

Given two manifolds $\mathcal{M}_1$ and $\mathcal{M}_2$ sampled by two point clouds $\mathcal{P}_1 = \{\mathbf{x}_i\}_{i=1}^n$ and $\mathcal{P}_2 = \{\mathbf{y}_i\}_{i=1}^n$, and $\mathcal{Q}_1 \in \mathbb{R}^{n\times n}$ and $\mathcal{Q}_2 \in \mathbb{R}^{n\times n}$ as two given pairwise descriptors, e.g., pairwise geodesic distance, between two points in $\mathcal{P}_1$ and $\mathcal{P}_2$ respectively. The shape correspondence problem can be casted as the following QAP as discussed in Section 1.1.2:

$$\underset{\mathbb{P}\in\Pi_n}{\arg\min} \|\mathbb{P}\mathcal{Q}_1 - \mathcal{Q}_2\mathbb{P}\|^2 = \underset{\mathbb{P}\in\Pi_n}{\arg\min} < \mathbb{P}, \mathcal{Q}_2\mathbb{P}\mathcal{Q}_1 > \tag{2.1}$$

where $P \in \mathbb{R}^{n\times n}$ is a permutation matrix with binary $\{0,1\}$ elements and each row and column sum is 1. Usually, most literature [10, 11, 76] uses Frobenius norm in equation 2.1.

Since the QAP problem is NP-hard [64], it is common to relax the permutation matrix in (1.2) to a doubly stochastic matrix, where $D \in \mathcal{D}_n = \{D \in \mathbb{R}^{n\times n} |\ D\vec{1} = \vec{1}, D^T\vec{1} = \vec{1}, D_{ij} \geq 0\}$, in the shape registration context [2, 11, 34, 41]. The doubly stochastic matrix representation not only convexifies the original QAP (1.2) but also provides a more general probabilistic interpretation of the map. The correspondence map of a point is spread as a probability distribution on the target shape. It is particularly useful in situations where two shapes do not have the same number of points or are sampled differently; in other words, where one to one correspondence map does not have a proper physical interpretation.

However, there remain at least two major computational challenges to solve the relaxed QAP for correspondence problems between shapes of relatively large size. First, the usual choice of pairwise descriptors, such as pairwise distance [77], heat kernel [14], and wave kernel [5] are represented as dense matrices and so are the doubly stochastic matrix. It can pose a storage and memory issue when two shapes are of large size even before conducting any computation. E.g., to store a full 40,000 by 40,000 dense matrix with 16-bit accuracy, it will take 12 GB RAM. When dealing with original relaxed QAP of size 40,000 by 40,000, there are three such matrices to be stored, which will cost 36 GB RAM in total.

In this case, certain approximation has to be used, such as sub-sampling methods which need remesh process and loses resolution [70], truncation of pairwise descriptors or spectrum approximation [3] which loses fine details in the shape, usually hard to decide what is the approximate spectral dimension, and may lead to accuracy problems due to the approximation error. Computationally, single dense matrix multiplication of the pairwise descriptor matrix and the doubly stochastic matrix requires $O(n^3)$ operations, where n is the number of points. When the number of vertices exceeds 10,000, the cost of one single matrix multiplication computation is already huge. More seriously, the relaxed QAP is usually solved by an iterative method. Due to the coupling of all elements of the doubly stochastic matrix, i.e., every element is affected by all other elements, elements corresponding to good matching can be influenced by those corresponding to the wrong matching initially, which can cause a slow convergence of the optimization process, especially when the initial guess is not good enough. Furthermore, for data with noise or distortion, the QAP may propagate the distortion or noise in one region to other regions and cause the solution to the QAP unsatisfactory. Lastly, the permutation matrix recovered from the doubly stochastic matrix may be far from the true solution due to the relaxation alongside noise in real data.

To tackle the aforementioned challenges for the QAP, we propose the following relaxed quadratic assignment using sparse pairwise descriptors, introduce a criterion to detect anchor pairs with good correspondence, and develop an efficient iterative algorithm with sparsity control for the doubly stochastic matrix using those high-quality, dense anchor pairs. These anchor pairs are then used in the final post-processing step to construct the full correspondence.

As a summary, the key ideas of SEQA are as follows:

- Sparse pairwise descriptor
- Relaxed QAP with Dynamically sparsity control
- Post Processing

2.2 Sparse pairwise descriptors

Let $(\mathcal{M}, g)$ be a closed 2-dimensional Riemannian manifold, the LBO is defined as:

$$\Delta_{(\mathcal{M},g)}\psi = \frac{1}{\sqrt{G}}\sum_{i=1}^{2}\frac{\partial}{\partial x_i}(\sqrt{G}\sum_{j=1}^{2}g^{ij}\frac{\partial\psi}{\partial x_j}) \tag{2.2}$$

where g^{ij} is the inverse of g_{ij} and $G = det(g)$ [16]. LBO is an elliptic and self-adjoint operator intrinsically defined on the manifold; thus, it is invariant under isometric transformation. The LBO eigen-system satisfies:

$$\Delta_{(\mathcal{M},g)}\psi_i = -\lambda_i\psi_i, \quad \int_{\mathcal{M}}\psi_i\psi_j \mathrm{d}s = \delta_{ij} \tag{2.3}$$

and uniquely determines the underlying manifold up to isometry [7]. Spectral geometry is widely used in shape analysis [59, 42, 74, 67, 13, 38, 50, 39, 65].

In practice, $\mathcal{M}$ is discretized by a triangular mesh $T = \{\tau_\ell\}$ with vertices $V = \{\mathbf{x}_i\}_{i=1}^{n}$ connected by edges $E = \{e_{ij}\}$. For each edge e_{ij} connecting points p_i and p_j, we define the angles opposite E_{ij} as angles α_{ij} and β_{ij}. Denote the stiffness matrix as $\mathbb{S}$, given by [56, 58]

$$\mathbb{S}_{ij} = \begin{cases} -\frac{1}{2}[cot\alpha_{ij} + cot\beta_{ij}] & i \sim j \\ \sum_{k\sim i}\mathbb{S}(i,k) & i = j \end{cases} \tag{2.4}$$

where $\sim$ denotes the connectivity relation by an edge. The mass matrix $\mathbb{M}$ is given by

$$\mathbb{M}_{ij} = \begin{cases} \frac{|\tau_1|+|\tau_2|}{12} & i \sim j \\ \sum_{k\sim i}\mathbb{M}(i,k) & i = j \end{cases} \tag{2.5}$$

where $|\tau_1|$ and $|\tau_2|$ are the areas of the two triangles sharing the same edge ij. On the one hand, the eigen-system of LBO can be computed as $\mathbb{S}\psi = \lambda\mathbb{M}\psi$, which suggests $\mathbb{S}$ and

$\mathbb{M}$ implicitly contain the spectrum information of LBO which can be used to determine a manifold uniquely up to isometry. On the other hand, it has been rigorously shown a global rigidity theorem on the Stiffness matrix, i.e. two polyhedral surfaces share the same Stiffness matrices if on only if their corresponding metrics are the same up to a scaling [27]. Note that the mass matrix fixes the scaling factor. Furthermore, both of these two matrices are local which are not sensitive to boundary conditions or topological perturbations. Additionally, to compute a Stiffness matrix and Mass matrix for a given mesh is very fast without involving solving any numerical PDE; most importantly, both of them are sparse which make them practical to be stored and computed with when vertices number is large.

Therefore, we expect that $\mathbb{S}$ and $\mathbb{M}$ together can serve as good sparse pairwise descriptors in a QAP formulation for shape correspondence.

2.3 Relaxed QAP with sparsity control

2.3.1 Relaxed QAP for shape correspondence

Given two surfaces $\mathcal{M}_1$ and $\mathcal{M}_2$ discretized by triangular meshes with vertices $\{\mathbf{x}_i\}_{i=1}^n$ and $\{\mathbf{y}_i\}_{i=1}^n$ respectively. We denote the corresponding stiffness matrices by $\mathbb{S}_1, \mathbb{S}_2$ and the corresponding mass matrices by $\mathbb{M}_1$ and $\mathbb{M}_2$. Representing a point-to-point mapping between $\mathcal{M}_1$ and $\mathcal{M}_2$ by a permutation matrix $P \in \Pi_n$, we propose the following QAP problem to construct the point-to-point mapping between these two surfaces:

$$\min_{P\in\Pi_n} \frac{1}{2}\|P\mathbb{S}_1 - \mathbb{S}_2 P\|_F^2 + \frac{\mu}{2}\|P\mathbb{M}_1 - \mathbb{M}_2 P\|_F^2, \tag{2.6}$$

Where μ is a balancing parameter. The stiffness matrix captures local geometric information, and the mass matrix captures local area information of the discretized surface. Both matrices have a sparsity structure with the number of nonzero entries linearly scaled with respect to the number of points. This nice sparse property of both matrices already alleviates the memory issue for large data sets significantly. Besides, since both descriptors only capture local geometric information, it potentially allows the proposed model to handle the partial matching problem, open surfaces, and topological changes.

Since the proposed QAP is NP-hard, we relax the permutation matrix to a doubly stochastic matrix representation $D \in \mathcal{D}_n = \{D \in \mathbb{R}^{n\times n} |\ D\vec{1} = \vec{1}, D^T\vec{1} = \vec{1}, D_{ij} \geq 0\}$ of the mapping:

$$\min_{D\in\mathcal{D}_n} \frac{1}{2}\|D\mathbb{S}_1 - \mathbb{S}_2 D\|_F^2 + \frac{\mu}{2}\|D\mathbb{M}_1 - \mathbb{M}_2 D\|_F^2 \tag{2.7}$$

As an advantage of this relaxation, each row of D can be interpreted as the probability of a point on $\mathcal{M}_1$ mapping to points on $\mathcal{M}_2$. Now the relaxed QAP (2.7) is convex and can be solved by well-known algorithms in convex programming. Here, we use projected gradient descent algorithm with Barzilai-Borwein step size solve this optimization problem (see details in Section 2.5).

2.3.2 Dynamically sparsity-enforced QAP

As we pointed out before, the relaxed QAP problem (2.7) is still difficult to solve if dense doubly stochastic matrices are used in the optimization process. To overcome those difficulties, we propose an iterative algorithm that:

- Selects candidates for well-matched pairs as anchor pairs.
- Enforces a dynamic sparsity structure of the doubly stochastic matrix by using the regularity of the map, i.e., nearby points on the source surface should be mapped to nearby points on the target surface, in the neighborhood of those paired anchor pairs in each iteration.

These two ingredients both reduce the computation cost in each iteration (only sparse matrices are involved) and increase the number of well-matched pairs quickly since only candidates for well-matched points are used to guide the iterations.

Local mapping distortion Test

To define a desired sparsity structure for the doubly stochastic matrix D in the relaxed QAP (2.7), we first need to detect candidates for well-matched pairs, or equivalently to remove those definitely ill-matched points, dynamically in each iteration. Motivated by the Gromov-Wasserstein distance [47] and the unsupervised learning loss in [29], we introduce the following criterion to quantify local mapping distortion of a mapping at a point on the source manifold.

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

$$\mathcal{F}_\gamma(\phi)(\mathbf{x}) = \frac{1}{|\mathcal{B}_\gamma(\mathbf{x})|} \int_{\mathbf{y} \in \mathcal{B}_\gamma(\mathbf{x})} DE_\phi(\mathbf{x}, \mathbf{y}) \mathrm{d}\mathbf{y} \tag{2.8}$$

where $DE_\phi(\mathbf{x}, \mathbf{y}) = \frac{1}{\gamma} |d_{\mathcal{M}_1}(\mathbf{x}, \mathbf{y}) - d_{\mathcal{M}_2}(\phi(\mathbf{x}), \phi(\mathbf{y}))|$ *is the difference between the geodesic distance* $d_{\mathcal{M}_1}, d_{\mathcal{M}_2}$ *on the two corresponding manifolds, and* $|\mathcal{B}_\gamma|$ *is the volume of* $\mathcal{B}_\gamma$*.*

We have the following straightforward properties:

1. If ϕ is an isometric map, $\mathcal{F}_\gamma(\phi)(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{M}_1, \gamma > 0$.
2. If $\mathcal{F}_\gamma(\phi)(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{M}_1$ for some $\gamma > 0$, ϕ is isometric.

In discrete setting, $\mathcal{M}_1$ is represented as $\{\mathbf{x}_i\}_{i=1}^n$, $\mathcal{M}_2$ is represented as $\{\mathbf{y}_i\}_{i=1}^n$and the map ϕ is discretized as a one-to-one map between $\{\mathbf{x}_i\}_{i=1}^n$ and $\{\mathbf{y}_i\}_{i=1}^n$. We use the following discrete approximation:

$$\mathcal{F}_\gamma(\phi)(\mathbf{x}_i) \approx \frac{\sum_{\mathbf{x}_j \in \mathcal{B}_\gamma(\mathbf{x}_i), \mathbf{x}_j \neq \mathbf{x}_i} \mathbb{M}_1(j,j) DE_\phi(\mathbf{x}_i, \mathbf{x}_j)}{\left(\sum_{\mathbf{x}_j \in \mathcal{B}_\gamma(\mathbf{x}_i), \mathbf{x}_j \neq \mathbf{x}_i} \mathbb{M}_1(j,j)\right)} \tag{2.9}$$

to quantify how much ϕ is distorted locally and use it to prune out those points that have large local mapping distortion in the next iteration for the QAP. Given the ground truth map, local mapping distortion is very small at most regions except for area where non-isometric distortion is large as shown in Figure 2.1. In practice, we specify a number ϵ and view x_i as a candidate of well-matched anchor pairs for ϕ if $\mathcal{F}(\phi)(\mathbf{x}_i) < \epsilon$. Together with $\phi(\mathbf{x}_i)$, we extract a collection of anchor pairs $\{(\mathbf{x}_i, \phi(\mathbf{x}_i))\}_{i=1}^k$ which are used to define sparsity pattern in the doubly stochastic matrix D dynamically in the relaxed QAP (2.7). It is important to note that current anchor pairs will be re-evaluated and updated in later iterations.

Remark: The computation complexity of local mapping distortion is $O(n)$ as we only need to check a fixed radius neighborhood for each point. Moreover, by normalizing the distance according to the local element area, the local mapping distortion is robust with respect to global scaling as well as local sampling variation. Figure 2.1 shows that, for the ground truth correspondence, local mapping distortion is very small almost everywhere except in regions where non-isometric distortion is large. Hence the proposed local mapping distortion may also serve as a good unsupervised error metric when ground truth is not available.

Dynamic Sparsity for Doubly Stochastic Matrices

Suppose a collection of anchor pairs $\{(\mathbf{x}_i, \phi(\mathbf{x}_i))\}_{i=1}^k$ have been selected using the local mapping distortion test. In the next iteration, a sub-QAP only involving points in the neighborhood of selected anchor pairs are solved. We further enforce a sparsity structure on the doubly stochastic matrix for the sub-QAP based on the following two rules.

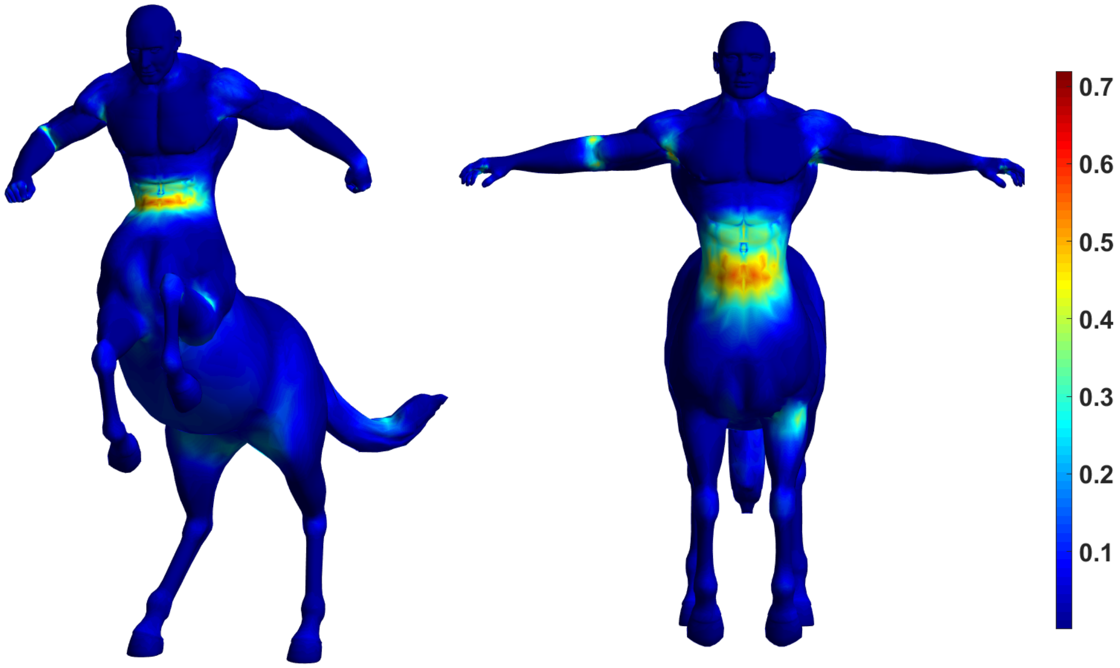

Figure 2.1: Ground Truth local mapping distortion: According to the color map, it is obvious that the value of local mapping distortion is large at the places where non-isometric distortion is large.

1. Each anchor pairs is mapped to its corresponding anchor pairs;
2. Points in the neighborhood of an anchor pair on the source surface are mapped to the neighborhood of the corresponding anchor pairs on the target surface.

Let $\mathcal{N}(\mathbf{x})$ denote the neighborhood of a given point on a manifold, e.g., a geodesic ball $\mathcal{B}_r$ centered at $\mathbf{x}$ on the manifold, or simply points in the l-th ring of $\mathbf{x}$ on a triangular mesh. Define $\mathcal{N}(\{\mathbf{x}_i\}_{i=1}^k) = \bigcup_{i=1}^k \mathcal{N}(\mathbf{x}_i)$ and $\mathcal{N}(\{\phi(\mathbf{x}_i)\}_{i=1}^k) = \bigcup_{i=1}^k \mathcal{N}(\phi(\mathbf{x}_i))$. For the doubly stochastic matrix D in the relaxed QAP (2.7), we only update variables with indices corresponding to the set $\mathcal{N}(\{\mathbf{x}_i\}_{i=1}^k) \times \mathcal{N}(\{\phi(\mathbf{x}_i)\}_{i=1}^k)$ together with the following sparsity constraints

$$D_{t,s} = \begin{cases} \delta_{\phi(\mathbf{x}_s),\mathbf{y}_t}, & \text{if } \mathbf{x}_s \in \{\mathbf{x}_i\}_{i=1}^{k} \\ 0, & \text{if } \mathbf{x}_s \in \mathcal{N}(\mathbf{x}_i) \text{ and } \mathbf{y}_t \notin \mathcal{N}(\phi(\mathbf{x}_i)) \\ 0, & \text{if } \mathbf{y}_t \in \mathcal{N}(\phi(\mathbf{x}_i)) \text{ and } \mathbf{x}_s \notin \mathcal{N}(\mathbf{x}_i) \end{cases} \cdot \tag{2.10}$$

By limiting the optimization region and enforcing the previous two sparsity constraints, the number of variables in the QAP problem after the sparsity enforcement is significantly reduced from $O(n^2)$ to $O(n)$. This can dramatically reduce computation cost. Moreover, since the anchor pairs are fixed in one iteration, it will no longer be influenced by other points in the current optimization process; on the contrary, it will enforce a positive influence on the neighboring points and rest points in a propagation sense.

In practice, we always choose the size of $\mathcal{B}_\gamma(\mathbf{x})$ in the distortion test smaller than the size of sparsity control neighborhood $\mathcal{N}(\mathbf{x})$ to allow the growth of anchor pairs in the next iteration. In our experiments, we choose $\mathcal{B}_\gamma(\mathbf{x}_i)$ as points included in the second ring of $\mathbf{x}_i$ and $\mathcal{N}(\mathbf{x}_i)$ as points included in the fourth ring of $\mathbf{x}_i$. The larger $\mathcal{B}_\gamma$ is, the more precise anchor pairs will be; the larger sparsity neighborhood $\mathcal{N}(\mathbf{x})$ is, the faster the number of anchor pairs grows. However, computation cost also increases for each QAP iteration when $\mathcal{B}_\gamma(\mathbf{x})$ and $\mathcal{N}(\mathbf{x})$ become larger since the doubly stochastic matrix is less sparse.

Once the sparsity regularized D is obtained, we update the point-to-point mapping ϕ by choosing the largest element in each row. Then, we find a new collection of anchor pairs by the distortion test based on the updated ϕ. Figure 2.2 illustrates an example of this procedure in the first 5 iterations.

Ideally, one should grow anchor pairs until all points are covered. However, because of noise and/or non-isometry, the growth of high-quality anchor pairs usually slows down after a

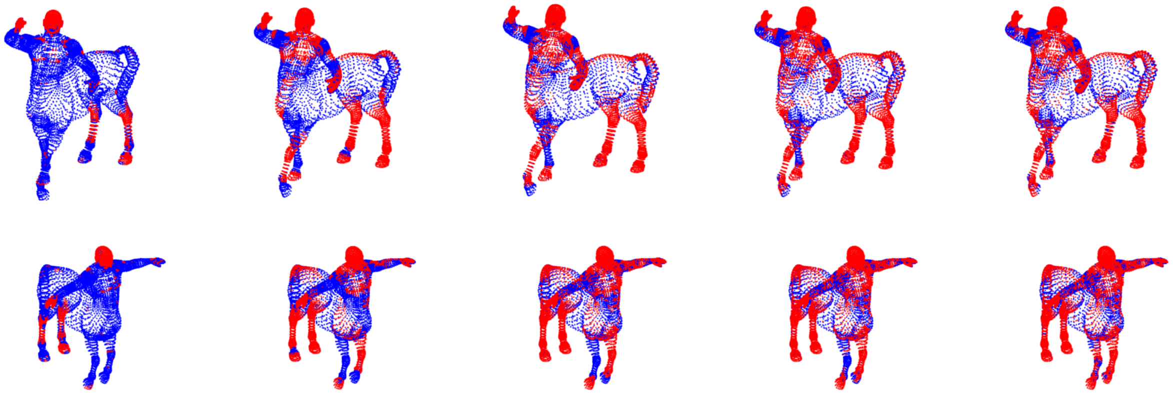

Figure 2.2: Example of the growth of anchor pairs. Red points in each column from left to right represent the location of anchor pairs selected by local mapping distortion test from iteration 1 to 5.

few iterations. Two point clouds describing two isometric manifold may not have exactly the same stiffness matrix and mass matrix up to a permutation. Moreover, even the exact solution of QAP (1.2) may not produce a desirable result. For some shape with a relatively large noise, the growth of anchor pairs may stop before covering all points. To balance between efficiency and accuracy, we find that five iterations of relaxed QAP (2.7) is good enough to find enough high-quality anchor pairs. Then a post-processing step is used to construct the correspondence for the remaining points with the help of matched anchor pairs. In some sense, we exhaust the usage of local feature (our sparse pairwise descriptor) first and then turn to global feature (post-processing). The main belief is that only relying on features under the same scale, either local or global, is not good enough for nearly isometric shape correspondence. A robust and accurate scheme should combine the benefits from both local and global features.

2.4 Post Processing

The aforementioned sparsity enforced quadratic assignment model based on local features is effective in growing anchor pairs from initial seed points in the regions where there is no significant non-isometric distortion. However, for regions where there is significant distortion, such as near fingertips or elbow regions of humans in different poses, local features may not be enough to produce satisfactory results.

To construct a full correspondence and improve mappings in those regions, one can use global pointwise descriptors with reference to those already well-matched anchor pairs. There are various options for the final post processing step. For example, we use Heat Kernel Signature (HKS) [67] for closed surfaces in our experiments.

To use HKS as a point-wise descriptor, we first need to compute the heat kernel. Given a compact manifold $\mathcal{M}_1$, the heat kernel are fundamental solutions to the heat diffusion,

$$\frac{\partial u(t,x)}{\partial t} = \Delta_{\mathcal{M}_1} u(t,x),\ with\ u(0,x) = u_0(x) \tag{2.11}$$

where $u : [0, +\inf) \times \mathcal{M}_1 \to \mathbb{R}$. The solution of heat diffusion is given by

$$u(t,x) = \int_{\mathcal{M}_1} k(t,x,y) u_0(y) dy,\ k(t,x,y) = \sum_i e^{\lambda_i t} \phi_i(x) \phi_i(y) \tag{2.12}$$

where ϕ and λ are eigen-functions and eigen-values of Laplacian operator. If we apply the heat kernel on delta functions at point p_i with time t, it is equivalent to measuring heat distribution after time t with a single heat source at point p_i. Hence HKS can provide pointwise heat feature obtained from each landmark, as shown in Figure 2.3. The choice of time parameter t will influence the quality of the mapping. Smaller t emphasizes the relationship between nearby heat source while larger t involves more information generated by far-away heat sources. Since the landmarks we pick are likely to be exact, so here we

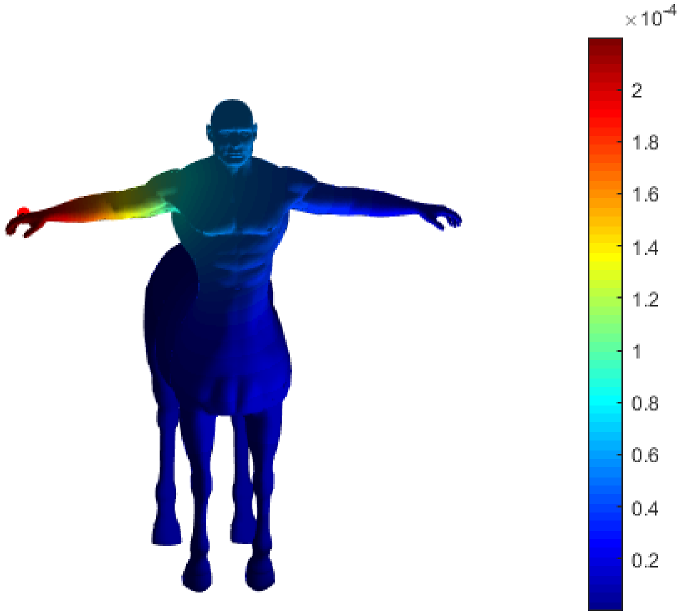

Figure 2.3: Example of the diffusion of heat from a source point. Red point denotes the heat source and color indicates the temperature at each point after some time.

always choose a relatively large t, such that all the heat sources can have a certain influence on the leftover points.

Let $H_1(\mathbf{x}, \mathbf{x}', t)$ and $H_2(\mathbf{y}, \mathbf{y}', t)$ denote HKS on $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively. Given $\{\mathbf{x}_i, \phi(\mathbf{x}_i)\}_{i=1}^{\ell}$ as the (sub)set of high quality anchor pairs obtained from solving the QAP (2.7), we construct pointwise descriptors of length ℓ for those $\mathbf{x} \in \mathcal{M}_1, \mathbf{y} \in \mathcal{M}_2$ not in the anchor pair set as $\{H_1(x, \mathbf{x}_i, t)\}_{i=1}^{\ell}, \{H_2(y, \phi(\mathbf{x}_i), t)\}_{i=1}^{\ell}$. Then we simply perform a nearest neighborhood search in this descriptor space to find the correspondence for non-anchor pairs.

Alternatively, other global pointwise features can also be used depending on different application scenarios. For example, for patches, we use geodesic distance to the chosen anchor pairs on the corresponding surfaces as the pointwise descriptor since it is more stable to the artificial boundary and topological changes.

2.5 Numerical Algorithms

We use projected gradient descent with Barzilai-Borwein step size [6], summarized in Algorithm 2.1, with the dynamic sparsity constraints (2.10) in each iteration to solve (2.7). The initial doubly stochastic matrix D_0 can be a random matrix or using the initial guess provided by SHOT feature [71] satisfying the sparsity constraint by projection (2.19).

The SHOT feature only needs to be computed once at the beginning to provide the initial doubly stochastic matrix D_0 and select anchor pairs for the first iteration. In later iterations, initial guess can be provided by projecting D from the previous iteration according to the new sparsity constraint.

Algorithm 2.1 Projected gradient decent for (2.7)

```
repeat
    1.$Y_{k+1} = D_k - \alpha_k \nabla_D (\|D_k \mathbb{S}_1 - \mathbb{S}_2 D_k\|_F^2 + \mu \|D_k \mathbb{M}_1 - \mathbb{M}_2 D_k\|_F^2)$
    2.$D_{k+1} = \arg\min_{D \in \mathcal{D}_n} \|D - Y_{k+1}\|_F^2$
until
```

Remark. The uniqueness of the projection step in Algorithm 2.1 can be proved as follows:

Proof. We want to show that there exists a unique minimizer for the problem

$$min_P \|P - T\|_F^2, \ s.t\ P\vec{1} = \vec{1}, P^T\vec{1} = \vec{1}, \ P \in \mathbb{R}^{n \times n} \tag{2.13}$$

The Lagrangian of this problem is

$$\mathcal{L} = \|P - \mathcal{T}\|_F^2 + \lambda < P\vec{1} - \vec{1}, \vec{1} > + \mu < P^T\vec{1} - \vec{1}, \vec{1} > \tag{2.14}$$

Then the KKT condition will be

$$\begin{cases} P - T + \lambda \cdot \vec{1}^{T} + 1 \cdot \mu^{T} \\ P\vec{1} = \vec{1} \\ P^{T}\vec{1} = \vec{1} \end{cases} \tag{2.15}$$

This system is underdertermined, there exist infinite solutions as long as they satisfies $n - |T| + n|\lambda| + n|\mu| = 0$. However we may assume that $|\lambda| = \frac{a}{n}(|T| - n)$ and $|\mu| = \frac{b}{n}(|T| - n)$, where $a + b = 1$. Hence P is uniquely defined by

$$P = T - \frac{1}{n}(T \cdot \vec{1} - \vec{1}) \cdot \vec{1}^{T} - \frac{1}{n}\vec{1} \cdot (T^{T} \cdot \vec{1} - \vec{1})^{T} + \frac{1}{n^2}(|T| - n) \cdot \vec{1} \cdot \vec{1}^{T} \tag{2.16}$$

□

Note that we only update entries of D corresponding to those points in the neighborhood of selected anchor pairs $\mathcal{N}(\{\mathbf{x}_i\}_{i=1}^{k}) \times \mathcal{N}(\{\phi(\mathbf{x}_i)\}_{i=1}^{k})$ and perform the projection on the set of doubly stochastic matrix D satisfying the sparsity constraint (2.10). Let $\mathcal{C}$ be the indicator matrix for the sparsity constraint

$$\mathcal{C}_{t,s} = \begin{cases} \delta_{\phi(\mathbf{x}_s),\mathbf{y}_t}, & \text{if } \mathbf{x}_s \in \{\mathbf{x}_i\}_{i=1}^{k} \\ 1, & \text{if } \mathbf{x}_s \in \mathcal{N}(\mathbf{x}_i) \text{ and } \mathbf{y}_t \in \mathcal{N}(\phi(\mathbf{x}_i)) \\ 0, & \text{otherwise} \end{cases} \cdot \tag{2.17}$$

The solution to the projection step in Algorithm 2.1

$$D_{k+1} = \operatorname*{arg\,min}_{D \in \mathcal{D}_n} \|D - Y\|_F^2, \text{ s.t. } (2.10) \tag{2.18}$$

is given by

$$D_{k+1} = \Big(Y + \frac{|Y_{\mathcal{C}}| - |\mathcal{C}|}{|\mathcal{C}|^2}\vec{1}\vec{1}^T - (Y_{\mathcal{C}}^T\vec{1} - \vec{1}) \oslash \vec{c}\vec{1}^T - \vec{1}((Y_{\mathcal{C}}\vec{1} - \vec{1}) \oslash \vec{r})^T\Big)_{\mathcal{C}} \tag{2.19}$$

where $(\cdot)_{\mathcal{C}} = (\cdot) \odot C$, $|\cdot| = \vec{1}^T(\cdot)\vec{1}$, $\mathcal{C}\vec{1} = \vec{r}$, $\mathcal{C}^T\vec{1} = \vec{c}$ and $\odot$, $\oslash$ are Hadamard product (element-wise product) and Hadamard division (element-wise division). We further relax our problem by neglecting the nonegative constraint as suggested in [1]. This strategy further reduces the computation cost without causing any problem in all of our experiments.

Our iterative method for the relaxed QAP (2.7) is summarized in Algorithm 2.2. Starting from an initial point-to-point map ϕ^0 (or converted from an initial doubly stochastic matrix), the three steps in each iteration are:

- Update the set of anchor pairs using (2.9);
- Update the doubly stochastic matrix by Algorithm 2.1 with sparsity constraint based on updated anchor pairs;
- Convert the doubly stochastic matrix to an updated point-to-point map by choosing the index of the largest element in each row.

Note that all anchor pairs are updated and improved (by decreasing local mapping distortion tolerance ϵ) during the iterations. Geometrically, our iterative method is like matching by region growing from anchor pairs. The local mapping distortion criterion allows us to efficiently and robustly select a few reasonably good anchor pairs from a fast process (but not necessarily accurate dense correspondence), such as SHOT. Then anchor pairs will grow as well as improve due to gradually diminishing local mapping distortion tolerance during iterations. In our experiments, we find enough high-quality anchor pairs after five iterations by decreasing ϵ from 5 to 1. Then we use these anchor pairs to construct the correspondence of remaining points in the final post-processing step as described in Section 2.4.

Algorithm 2.2 Iterative method for relaxed QAP with dynamic sparsity control

```
**Input** a point-to-point map $\phi^0$, iteration steps n, $\{\epsilon_i\}_1^n$ and parameter $\mu$.
**repeat**
    1. Find anchor pairs $\{(\mathbf{x}_i, \phi^k(\mathbf{x}_i)) \mid \mathcal{F}(\phi^k)(\mathbf{x}_i) < \epsilon_k\}$. Define $\mathcal{N}_1^k = \mathcal{N}(\{\mathbf{x}_i\}_{i=1}^k)$ and $\mathcal{N}_2^k = \mathcal{N}(\{\phi^k(\mathbf{x}_i)\}_{i=1}^k)$.
    2. Compute $D^{k+1}$ by Algorithm 2.1 with sparsity constraint (2.10) on $\mathcal{N}_1^k \times \mathcal{N}_2^k$.
    3. Update $\phi^{k+1}(x_s) = y_t$, where $t = \arg\max D^{k+1}(s,:)$.
**until** n steps are reached
```

Since we start with a relatively large local distortion tolerance for initial anchor pairs, our method is quite stable with respect to the initialization. Moreover, as we decrease the tolerance with iterations, anchor pairs selected earlier can be updated in later iterations. We remark that the above algorithm can be straightforwardly extended to shape correspondence between two point clouds with different sizes by using a rectangular doubly stochastic matrix with the right dimension.

2.6 Discussion

2.6.1 Point cloud matching

We can easily extend our method to point clouds, raw data in many applications, without a global triangulation by constructing the stiffness and mass matrices at each point using the local mesh method [37] with an adaptive-KNN algorithm.

In [37], the local connectivity of a point p on the manifold M is established by constructing a standard Delaunay triangulation in the tangent plane at p of the projections of its K nearest neighbors. However, the classical KNN with fixed K is not adaptive to local geometric feature size or sampling resolution, which may lead to a loss of accuracy, as shown in Figure 2.4. So we introduce an adaptive-KNN.

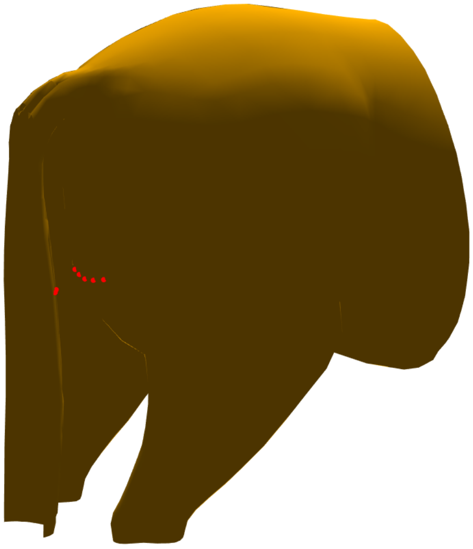

Figure 2.4: Incorrect connectivity caused by traditional KNN. Red points should not be connected. But traditional KNN will connect them because points are close in the embedding space, which will cause error in later diffusion process.

For neighborhood $\mathcal{N}(i)$ of p_i formed by traditional KNN, the local normalized covariance matrix $\mathbb{P}$ of this patch is defined as:

$$\mathbb{P} = \sum_{p_j \in \mathcal{N}(i)} (p_j - c)^T (p_i - c) \tag{2.20}$$

Where $c = \frac{1}{|\mathcal{N}(i)|} \sum_{p_j \in \mathcal{N}(i)} p_j$ is the local barycenter. The eigen-vector (e_1^i, e_2^i, e_3^i) of $\mathbb{P}$ forms an orthogonal local frame of the patch, and corresponding eigen-value $(\lambda_1^i, \lambda_2^i, \lambda_3^i)$, $\lambda_1^i \geq \lambda_2^i \geq \lambda_3^i$, reflects the sum of the projections of all centralized points in each eigen-vector direction normalized by number of points. Since we want our local patch to be as flat as possible, the local patch should be constrained mostly in a two-dimensional manifold which means the last eigen-value should not be too large. In order to get rid of the influence the area of different patches, we further normalize λ_3^i by λ_1^i. Specifically, we set a threshold on $R = \frac{\lambda_3^i}{\lambda_1^i}$ which is invariant of local density and patch area to identify flat and smooth local patches.

The algorithm starts with a large K and gradually shrinks the size of K until the ratio R is smaller than a certain threshold of r, as shown in Algorithm 2.3. Our Adaptive KNN performs better than traditional KNN in the diffusion process, as shown in Figure 2.5.

Algorithm 2.3 Adaptive KNN

```
Given maximum neighborhood size $K$, eigen-value ratio threshold $r$ and shrink size $m$, for point $p_i \in P$ with initial neighborhood points set $\mathcal{N}_0(i)$ of size $K$
Initialization:
1.Compute local covariance matrix $\mathbb{P}_0$
2.Comute $R = \frac{\lambda_{min}(P_0)}{\lambda_{max}(P_0)}$
while $R \geq r$ do
    1.Update $\mathcal{N}_{k+1}(i)$ by excluding furthest $m$ points from $\mathcal{N}_k(i)$
    2.Update local covariance matrix $\mathbb{P}_{k+1}$
    3.Update $R = \frac{\lambda_{min}(P_{k+1})}{\lambda_{max}(P_{k+1})}$
end while
```

(a) Heat Diffusion using traditional KNN

(b) Heat Diffusion using adaptive KNN

Figure 2.5: Heat diffusion at tail part (indicated in red circle) is improved a lot by using our adaptive KNN to construct the mesh. Red dot represents the heat source and warmer color with higher temperature.

2.6.2 Patch/partial matching

In real applications, well-sampled data for 3D shapes are not easy to obtain. Instead, holes, patches, or partial shapes are more common in real data. Correspondence between shapes with topological perturbations, artificial boundaries, and different sizes are difficult for methods based on global intrinsic descriptors in general. For example, the spectrum of LBO is sensitive to boundary conditions and topological changes.

However, since our method is based on local features, the effect of boundary conditions and topological perturbations is localized. Hence our method can be directly applied to those scenarios with excellent performance. For example, our iterative method for the relaxed QAP using anchor pairs and sparsity control fits the smaller patch into the larger one nicely for partial matching (see Figure 1.3). For post processing in patch matching, we switch from HKS to geodesic distance signature since HKS is sensitive to boundary conditions.

2.7 Experiment Results

We evaluate the performance of our method through various tests on data sets from TOSCA [12] and SCAPE [4] and on patches extracted from TOSCA. All inputs for our tests are raw data without any preprocessing, i.e., no low-resolution model or pre-computed geodesic distance. Experiments are conducted in Matlab on a PC with 16GB RAM and Intel i7-6800k CPU. The result of our method using mesh input is denoted as the mesh method, and the result of our method without using mesh is denoted as point cloud method.

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

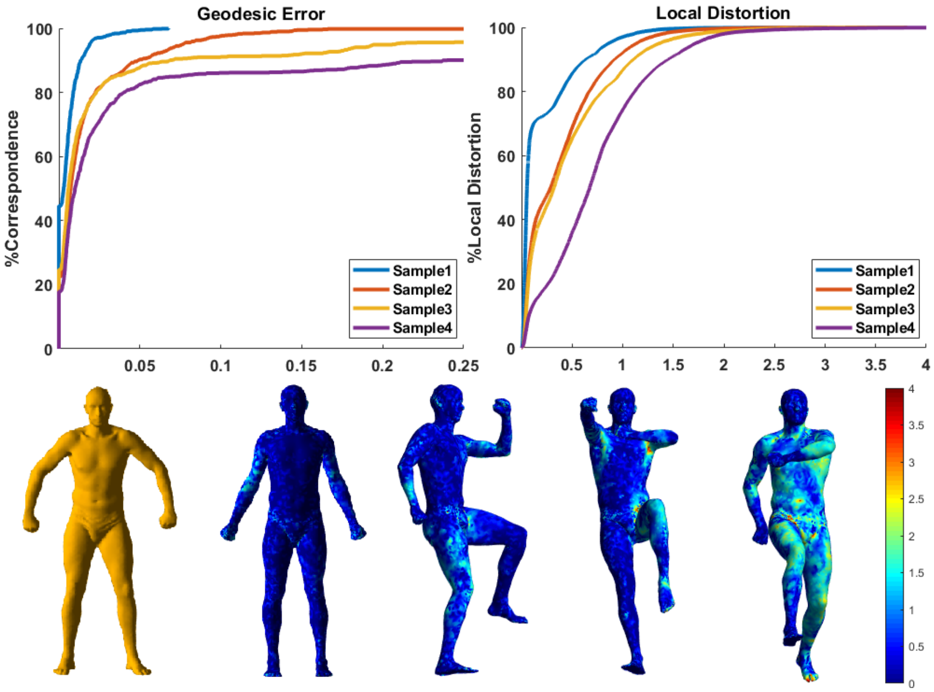

Figure 2.6: Geodesic error and cumulative density function of local mapping distortion for 4 sample pairs from the SCAPE data set (1st row); Original shape and corresponding local mapping distortion heat map for sample 1 to 4 (2nd row).

local mapping distortion defined in (2.8) can also serve as an unsupervised error metric to measure the quality of a map. As shown in Figure 2.6, it's clear that local mapping distortion decreases as the geodesic error decreases, which indicates that local mapping distortion can serve as a good unsupervised metric to quantify the approximate isometry.

TOSCA The TOSCA data set contains 76 shapes of 8 different classes, from humans to animals. The number of vertices varies from 4k to 50k. We use 5 iterations to grow the set of anchor pairs. The neighborhood used for local mapping distortion test for selecting anchor pairs is the second ring, and for sparsity control is the fourth ring. The distortion threshold decreases equally during the iterations from 5 to 1; the gradient descent step size in Algorithm 2.1 is 75; we approximate the heat kernel by 300 eigen-functions of the LBO with a diffusion time $t = 50$ in the post-processing step. For point clouds without mesh, we use an initial $K = 200$, ratio $r = 0.05$, and shrink size $m = 6$ for our adaptive-KNN; HKS post processing is not used since the spectrum computed directly from the point cloud

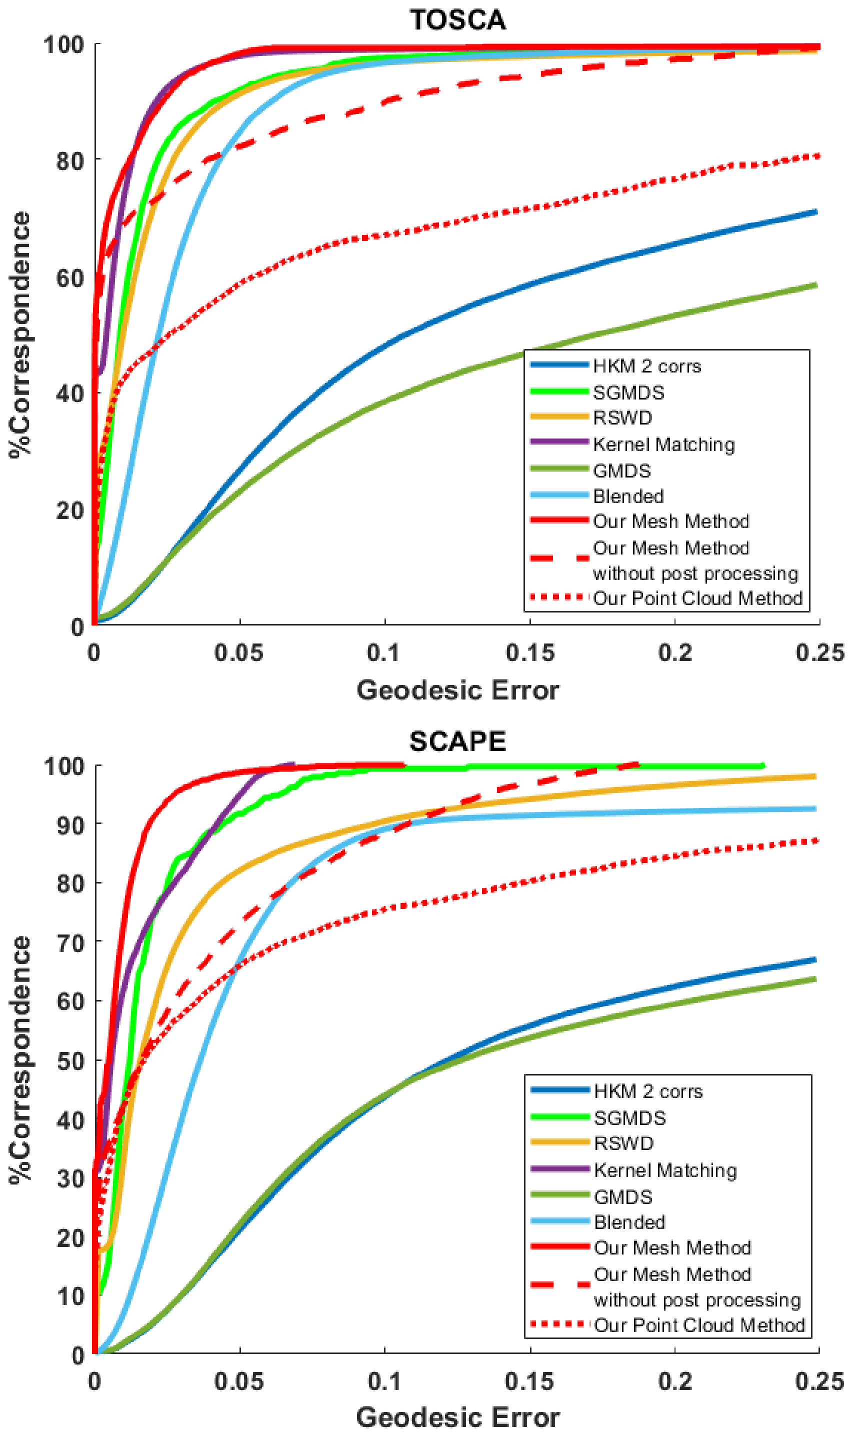

Figure 2.7: Correspondence accuracy on the TOSCA and SCAPE data sets.

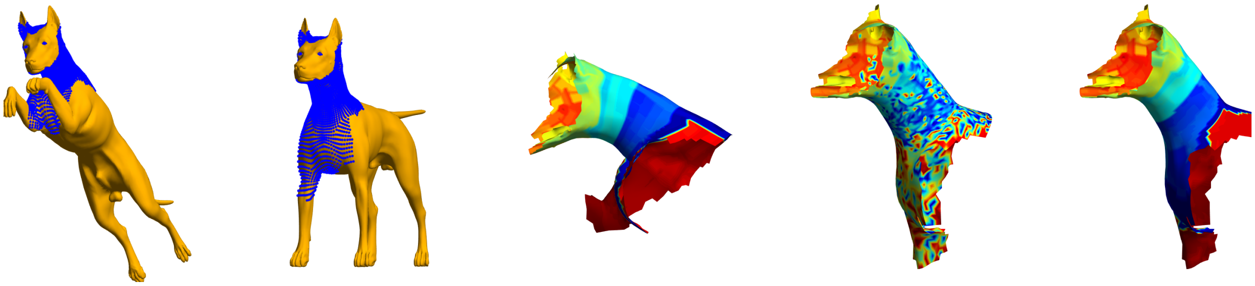

Figure 2.8: An example of matching two patches. The first two images show the location of the patches; the third image is the color map on the first patch; the fourth image is the induced color map based on SHOT feature; the last image is the induced color map from our mesh method.

is not accurate enough. Results of our mesh method with or without post processing, and point cloud method without post processing are presented. We compare our method with the following methods: Blended [36], SGMDS [3], GMDS [11], Kernel Marching [76], RSWD [39], and HKM 2 corrs [51].

Figure 2.7 shows the quantitative result in terms of the geodesic error metric. Our mesh method outperforms most of the state-of-art methods. Our mesh method without post-processing and point cloud method also achieve reasonably good results.

SCAPE The SCAPE data set contains 72 shapes of humans in different poses. Each shape has 12,500 vertices. We use the same parameters as those on TOSCA data set except for diffusion time $t = 0.001$ in the post-processing step. Results of our mesh method with or without post processing, and point cloud method are presented.

Figure 2.7 shows the quantitative result. Our method achieves state-of-art accuracy. Again, our mesh method without post-processing and point cloud method also achieve reasonably good results.

Patch Matching We present a few test results for patches that have holes, boundaries, and partial matching. We paint the first patch with colors and map the color to the second patch with the correspondence computed using SHOT [71] as the pointwise descriptor, which

Figure 2.9: Examples of matching patches with topological perturbation and shapes with randomly missing elements. The first and third columns illustrate the patches and shapes to match. The top color map of the first patch/shape is mapped to the second patch/shape using SHOT (middle) and our method (bottom).

also serves as the initial guess for our method, and the correspondence computed from our algorithm to visualize the result. Since HKS is sensitive to boundary conditions and topological changes, we use geodesic distance to those selected anchor pairs as pointwise descriptors in the post processing step.

The first test is matching two patches of a dog with different poses from TOSCA, as shown in Figure 2.8. The two patches have very irregular boundaries. Using extrinsic pointwise descriptors, such as SHOT, fail to give a good dense correspondence. However, our method performs well since it uses local pairwise descriptors to find high-quality anchor pairs and integrates global pointwise descriptors, the geodesic distance to those anchor pairs, to complete the dense correspondence.

The second test is matching two patches with topological perturbations from TOSCA data, as shown in Figure 2.9. The first case is two different poses of a wolf with mesh elements

Model	Wolf	Centaur	Horse	Cat	David
Number of Vertices	4344	15768	19248	27894	52565
Mesh Method(s)	59	531	801	929	1681
Point Cloud Method(s)	57	524	811	937	1610

Table 2.1: Average run time for shapes from TOSCA

randomly deleted from each surface independently. The second case is body parts of a cat in different poses with topological perturbation; the second patch is not connected at the bottom while the first one is as highlighted in the figure. Since neither local connectivity distortion nor missing elements will significantly influence the stiffness matrix or mass matrix at most points, our method can still produce good results.

We further test our method on a pair of patches with both different sizes (partial matching) and topological changes, as shown in Figure 1.3. The example is mapping an arm patch to an entire shape. Extra points are colored as blue. The arm patch has fewer points than the entire shape, and the figure tips are cut off, which results in both different sizes and topological changes. Even for this challenging example, our method performs well.

Time efficiency We list the average run time of several examples in TOSCA data set in Table 2.1. Most state-of-the-art methods using (dense) pairwise descriptors and quadratic assignment (QA) require dense matrix multiplication in each step, which already has super-quadratic complexity. Although Laplace-Beltrami (LB) eigen-functions can be used to compress the dense matrix by low-rank approximation, it is still less sparse or localized and more time-consuming to compute than our simple, sparse and localized pairwise descriptor. Combined with our sparsity-enforced method for QA, our method has at most $O(n^2)$ complexity, which outperforms methods with super-quadratic complexity when handling data with large size. Experimentally, our method shows complexity even better than $O(n^2)$.

2.8 Conclusion

We develop a simple, effective iterative method to solve a relaxed quadratic assignment model through sparsity control for shape correspondence between two approximate isometric surfaces based on a novel local pairwise descriptor. Two key ideas of our iterative algorithm are:

1. select pairs with good correspondence as anchor pairs using a local unsupervised distortion test;
2. solve a regularized quadratic assignment problem only in the neighborhood of selected anchor pairs through sparsity control;
3. With enough high-quality anchor pairs, various pointwise global features with reference to these anchor pairs can further improve the dense shape correspondence.

Extensive experiments are conducted to show the efficiency, quality, and versatility of our method on large data sets, patches, and purely point cloud data. Similar to many existing methods, our method will have difficulty in dealing with significant non-isometric distortion and highly non-uniform sampling.

Chapter 3

Dual Iterative Refinement

In this chapter, we explain our *Dual Iterative Refinement* (DIR) method in detail. We present the idea and numerical algorithm of our proposed method DIR in Section 3.3 followed by a discussion about a few possible extensions of our method in Section 3.4. 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 3.5.

3.1 Functional map

Spectral geometry is widely used in shape analysis [59, 42, 74, 13, 38, 50, 39, 65, 30]. As an important member in spectral geometry, functional map is introduced in [50] to solve non-rigid shape correspondence problem. It improves earlier point-based spectral methods [31, 32, 45, 52] which directly matches the spectral embeddings of shapes. We first provide a brief introduction to the concept of functional map which is used in our multi-scale dual iterative refinement process.

Given $(\mathcal{M}_1, g)$ as a closed 2-dimensional Riemannian manifold, the LB operator uniquely determines the underlying manifold up to isometry [7]. 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 $\mathcal{L}_{\mathcal{M}_1} = \mathcal{A}_1^{-1}\mathcal{W}_1$ [55], 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 $\Phi^k_{\mathcal{M}_1} \in \mathbb{R}^{n\times k}$ whose rows are the first k LB eigen-functions evaluated at each point, $\Phi^k_{\mathcal{M}_1}(x) = (\phi_1^{\mathcal{M}_1}(x), ..., \phi_k^{\mathcal{M}_1}(x))$. Similarly, we can define the embedding for $\mathcal{M}_2$ as $\Phi^k_{\mathcal{M}_2} = \{(\phi_1^{\mathcal{M}_2}(x), ..., \phi_k^{\mathcal{M}_2}(x)) \mid x \in \mathcal{M}_2\}$.

The key idea of functional map is to have a linear representation of a shape correspondence $T : \mathcal{M}_1 \to \mathcal{M}_2$ in the functional spaces $C(\mathcal{M}_1)$ and $C(\mathcal{M}_2)$. In discretization, the functional map for a given correspondence T is a linear map (a matrix in discrete setting) C which aligns given bases of $\mathcal{C}(\mathcal{M}_1)$ and $\mathcal{C}(\mathcal{M}_2)$. In practice, LB embedding $\Phi_{\mathcal{M}}$ is commonly chosen as the basis for the functional space $\mathcal{C}(\mathcal{M})$. More precisely, 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:

$$C = \arg\min_{C} \|\Phi_{\mathcal{M}_1}C - \Pi\Phi_{\mathcal{M}_2}\|_F^2 = \Phi^{\dagger}_{\mathcal{M}_1}\Pi\Phi_{\mathcal{M}_2} \tag{3.1}$$

where † denotes the Moore-Penrose pseudo-inverse. Theoretically, it is well-known that LB eigen-functions are isometrically invariant. In other words, if $\mathcal{M}_1$ and $\mathcal{M}_2$ are isometric, then the corresponding LB eigen-functions are the same up to possible ambiguities, which form an orthonormal group. This motivates a constrained version of the following functional map [39]

$$C = \arg\min_{C\in\mathfrak{R}} \|\Phi_{\mathcal{M}_1}C - \Pi\Phi_{\mathcal{M}_2}\|_F^2 = \mathrm{Proj}_{\mathfrak{R}}(\Phi^{\top}_{\mathcal{M}_1}\Pi\Phi_{\mathcal{M}_2}) \tag{3.2}$$

where we write the set of orthonormal matrices as $\mathfrak{R} = \{C \mid C^{\top}C = Id\}$ and the projection to

$\mathfrak{R}$ is provided as $\mathrm{Proj}_{\mathfrak{R}}(A) = UV^\top$ with a singular value decomposition (SVD) of $A = U\Sigma V^\top$. Since this paper aims at tackling isometric or nearly isometric shape matching problem, we intend to use (3.2) which is the equivalent to (3.1) for the isometric case.

In the typical pipeline of an iterative functional map based shape correspondence algorithm, the functional map is obtained via a least square optimization with various constraints and regularizations, such as preservation of given landmarks, communitivity with LB operator, and sparsity. Once the optimized functional map C is computed, the new correspondence map T can be obtained by solving the following problem:

$$T(p) = \arg\min_{q\in\mathcal{M}_2} \|\Phi_M(p) - \Phi_N(q)C^\top\|_F^2, \quad \forall p \in \mathcal{M}_1 \tag{3.3}$$

However, the determination of an accurate functional map C can be challenging when limited prior knowledge or a poor initial correspondence is available. Moreover, LB eigenembeddings becomes less stable when high eigen-modes of LB operator is used. On the other hand, confining to lower LB eigenmodes limits the resolution of the mapping, which leads to inaccurate shape correspondence. A natural multi-scale idea is to start the functional map from a coarse resolution in spectral embedding involving a few low eigen modes and construct the correspondence T, then iteratively add more and more higher eigen-modes to refine the map. The main motivation is that the functional map on a coarse scale (a small matrix) can be determined efficiently and stably from an initial approximate or limited correspondence, which then provides an improved correspondence at the next iteration. This strategy leads to iteratively computing (3.2) and (3.3), or (3.1) and (3.3), which is the key idea proposed in [39] using rotation-invariant sliced Wasserstein distance and in the Zoom-out process [46]. As a crucial component of these iterative refinement methods, the functional map C is obtained from the correspondence information of all points in the previous iteration. This intermediate correspondence, especially at the beginning, can be quite inaccurate. Consequently, this may lead to errors when computing the functional map at the next level.

We provide both theoretical and experimental evidence on the effect of correspondence error on spectral alignment, i.e., the functional map C in (3.2), which can be problematic for any iterative refinement procedure based on spectral geometry.

Assume we have two perfectly isometric manifold $\mathcal{M}_1$ and $\mathcal{M}_2$ and their corresponding discrete spectral embedding $\Phi_{\mathcal{M}_1},\ \Phi_{\mathcal{M}_2} \in \mathbb{R}^{n\times k}$, where n is the number of points and k is spectral dimension. 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 between $\Phi_{\mathcal{M}_1}$ and $\Phi_{\mathcal{M}_2}$ is an orthonormal matrix $C_T \in \mathfrak{R}$, i.e. $\Phi_{\mathcal{M}_2} = \Phi_{\mathcal{M}_1} C_T$ and $C_T = \mathrm{Proj}_{\mathfrak{R}}(\Phi_{\mathcal{M}_1}^\top \Phi_{\mathcal{M}_2})$.

Theorem 1. *Given $\Phi_{\mathcal{M}_2} = \Phi_{\mathcal{M}_1} C_T$ with $C_T \in \mathfrak{R}$. Let's 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 rest part of Π inaccurately maps $\Phi_{\mathcal{M}_1}$. 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*

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

Then it is most likely that the spectral norm $\|C_a - C\|_2 > 0$ with probability at least $1-\eta$ where $\eta = \sum_{j=0}^{\lfloor n_2/2\rfloor} \frac{1}{2^j j!(n_2-2j)!}$

Notice that η decreases rapidly as n_2 grows. For example, when $n_2 = 25$, $\eta \approx 10^{-12}$.

Proof. We first show that a necessary condition for $\|C_a - C_T\|_2 = 0$ is that permutation σ is an involution, i.e. symmetric matrix.

Since $\Phi_{\mathcal{M}_2} = \Phi_{\mathcal{M}_1} C_T$, we have $Y_i = X_i C, i = 1, 2$. It is clear that $\Phi_{\mathcal{M}_1}^\top \Pi \Phi_{\mathcal{M}_2} = (X_1^\top X_1 + X_2^\top \sigma X_2) C_T$. Assume that there is a SVD decomposition $X_1^\top X_1 + X_2^\top \sigma X_2 = U \Sigma V^\top$. Then we have a SVD decomposition $\Phi_{\mathcal{M}_1}^\top \Pi \Phi_{\mathcal{M}_2} = U \Sigma V^\top C_T$ because C_T is an orthonormal matrix. Therefore, $C_a = \text{Proj}_{\mathfrak{R}}(\Phi_{\mathcal{M}_1}^\top \Pi \Phi_{\mathcal{M}_2}) = U V^\top C_T$. This yields

$$\|C_a - C_T\|_2 = \|UV^\top C_T - C_T\|_2 = \|UV^\top - \text{Id}\|_2$$

Thus, $\|C_a - C_T\|_2 = 0$ implies $U = V$. This leads to σ is symmetric. From the fact that there are $n_2!$ permutation matrix of size $n_2 \times n_2$ and there are $n_2! \sum_{j=0}^{\lfloor n_2/2 \rfloor} \frac{1}{2^j j! (n_2 - 2j)!}$ symmetric permutation matrix of size $n_2 \times n_2$ [66], it concludes the proof. □

One obvious observation is that $\eta = \sum_{j=0}^{\lfloor n_2/2 \rfloor} \frac{1}{2^j j! (n_2 - 2j)!}$ decreases rapidly as n_2 grows. For example, when $n_2 = 25$, $\eta \approx 10^{-12}$. However, to give a more quantitative characterization of the perturbation for an arbitrary shuffling is difficult since it depends not only on σ but also on X_1 and X_2. Instead, we conduct a few numerical experiments to demonstrate how an inaccurate correspondence Π will affect the correlation matrix $\Phi_{\mathcal{M}_1}^\top \Pi \Phi_{\mathcal{M}_2}$, which is the essential information to align spectral basis, e.g., functional map, between two shapes.

In our experiments, $\Phi_{\mathcal{M}_1}, \Phi_{\mathcal{M}_2} \in \mathbb{R}^{n \times k}$, where n is the total number of points and k is the spectral dimension. We map a human shape with 12,500 vertices to itself, which is a perfect isometric shape correspondence problem with the identity map as the ground truth. Theoretically, $0 \leq \|C_a - C_T\|_2 \leq 2$, since both C_a and C_T are orthonormal. The first experiment tests the behavior of $||\mathcal{C}_T - \mathcal{C}_a||_2$ with respect to the ratio of $\frac{k}{n_2}$ for two fixed spectral dimension, $k = 50$ and $k = 100$; second experiments tests the behavior of $||\mathcal{C}_T - \mathcal{C}_a||_2$ with respect to n_2 for two fixed $\frac{k}{n_2}$ ratio. A random permutation is imposed on Y_2 to compute C_a, and we independently run 50 trails for each parameter combination. Box-plots are used to illustrate the statistics of our computation in Figure 3.1. The experiments indicate that using inaccurate correspondences from the current level will most likely introduce error to the functional map for the next level, and the error can be quite significant (as big as the worst case, $\|C_a - C_T\|_2 = 2$). This could cause a failure or slow convergence for an iterative refinement strategy.

We further test a real example on one pair of shapes from TOSCA data set using ZoomOut [46], which is a simple iterative refinement method based on the functional map in the spectral domain. A fixed increment of spectral dimension is pre-specified for each iteration. Due to the use of current correspondence at all points to construct the functional map, the following iterative refinements may not lead to a satisfactory result in the end. In this test, we use 1000 LB eigen-functions for ZoomOut (and our method) with two different initial setups: (1) correspondence provided by SHOT, or (2) 4 given landmarks. For ZoomOut, we first compute a functional map for the first 4 spectral modes from the initial correspondence and then use the code provided by the authors on GitHub to run the experiment. Test results are plotted in Figure 3.2. These experiments further verify our analysis on possible issues using simple iterative spectral refinement without filtering out wrong correspondences.

Another critical problem for multi-scale approach in spectral domain is how to increase the resolution, i.e., spectral mode increment during the iterations, in an optimal, stable, and data-adaptive way. Most approaches [76, 39, 46] just adopt an empirical, prefixed, and inefficient increasing sequence.

Motivated from the above limitations of existing refinements solely in the spectral domain, we propose a dual iterative refinement method. The first key idea is to choose well-matched pairs from current correspondence, called anchor pairs, and only use them to determine the functional map accurately by which will, in turn, help to construct a much-improved correspondence to use in 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, which integrates local spatial information in the spectral refinement process. The second key idea is to find the maximal dimension of spectral modes that can

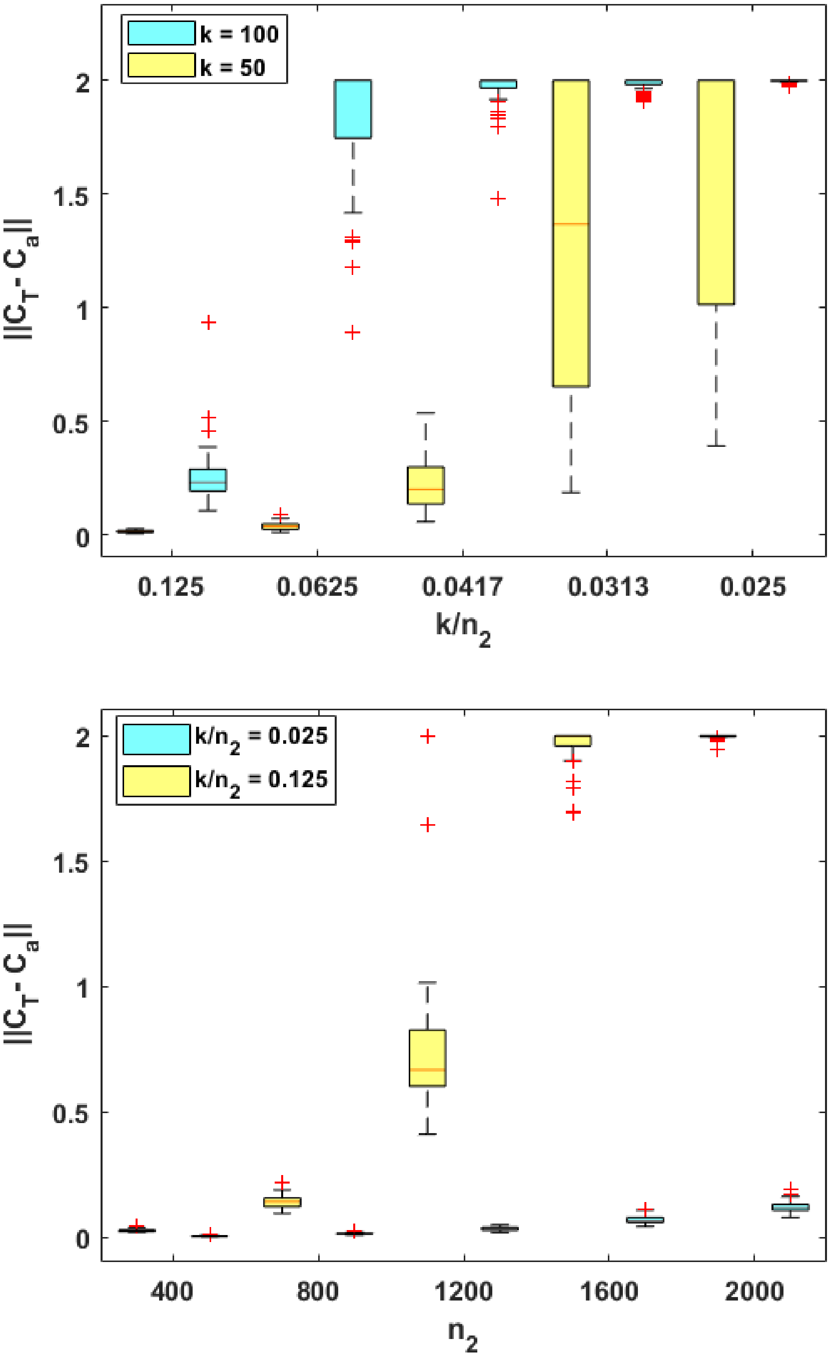

Figure 3.1: *Top:* Error statistics for the relation between functional map error and $\frac{k}{n_2}$. *Bottom:* Error statistics for the relation between functional map error and number of points in X_2.

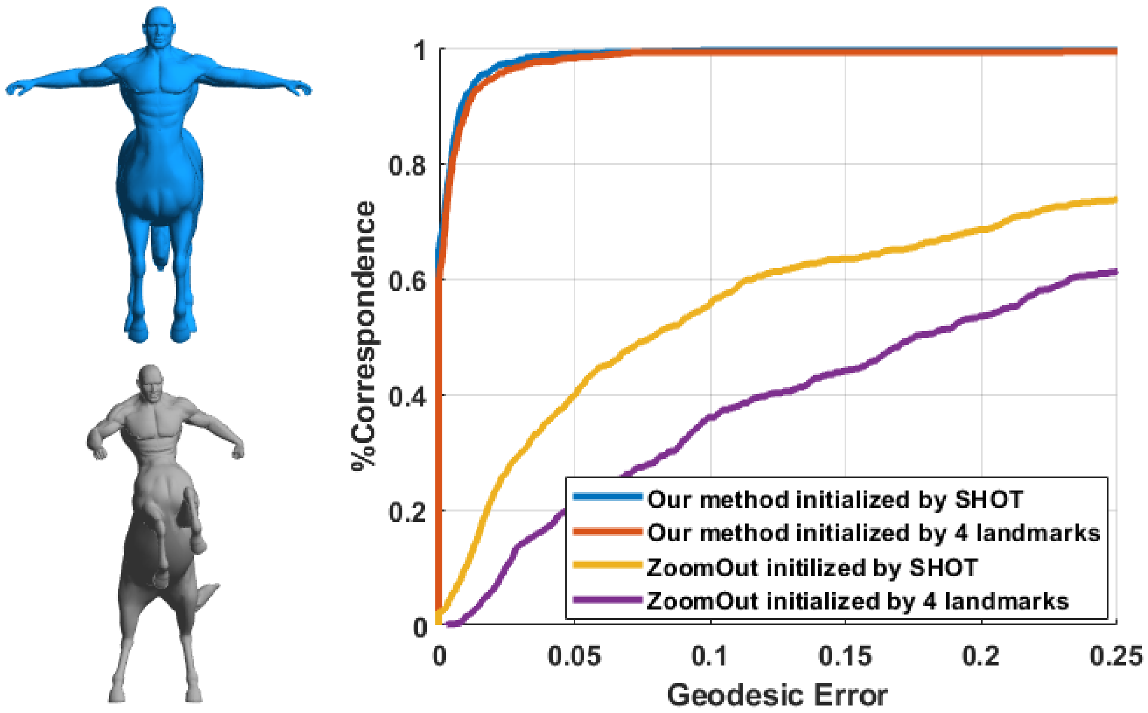

Figure 3.2: *Left:* The blue centaur is mapped to the gray centaur from TOSCA data set. *Right:* Geodesic error of our method and ZoomOut using different initial guess.

be robustly and accurately determined by these anchor pairs according to their distribution in spectral embedding using singular value analysis on their correlation matrix.

3.2 Local mapping distortion

In order to choose well-matched anchor pairs to compute the functional map C without any information of ground truth, we use the *local mapping distortion* introduced in Section 2.3.2 to measure the quality of correspondence for each pair and choose those pairs with small local mapping distortion.

3.3 Dual Iterative Refinement

We first precompute K eigen-functions for the discrete LB operator for each shape. K is determined by possible computation cost limitation, the noise level in the data, and the desired accuracy of the correspondence. In any case, it should not exceed what the mesh resolution can support for each 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 [71] 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 very accurate) methods. Then we start DIR which includes the following three simple steps, explained in detail later, in each iteration.

- Choose anchor pairs from current correspondence using local mapping distortion criterion;
- Determine the spectral dimension and the corresponding functional map based on the selected anchor pairs;
- Update the correspondence using the updated functional map.

We enforce two stopping criteria, whichever is satisfied first. The first is the total number of iterations, which limit the computation cost. The second is when the spectral dimension supported by anchor pairs reaches K. Here are more detailed descriptions of each step.

Step 1 Selecting anchor pairs.

By setting a proper local mapping distortion threshold ϵ, the set of pairs $\{(x_\ell, T(x_\ell)|\mathcal{D}_\gamma(T)(\mathbf{x}_\ell) < \epsilon\}$ are selected as anchor pairs from the current correspondence which will be used to guide next refinement. 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 and $\{T(x_\ell)\}_{\ell=1}^m \subset \mathcal{M}_2$, 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 spectral embedding, which can be quite smaller than the number of anchor pairs. We use singular value decomposition (SVD)[72], a well-known tool for dimension reduction, to find the dimension well expanded by a set of anchor pairs in the spectral domain. Consider the spectral embedding of a set of anchor pairs $\{x_\ell\}_{\ell=1}^m$ and $\{T(x_\ell)\}_{\ell=1}^m$ as $\Phi_{\mathcal{M}_1}^K(\{x_\ell\}_{\ell=1}^m)$, and $\Phi_{\mathcal{M}_2}^K(\{T(x_\ell)\}_{\ell=1}^m)$, respectively. We compute SVD of the correlation matrix of all anchor pairs

$$U\Sigma V^\top = \Phi_{\mathcal{M}_1}^K(\{x_\ell\}_{\ell=1}^m)^\top \Phi_{\mathcal{M}_2}^K(\{T(x_\ell)\}_{\ell=1}^m). \tag{3.4}$$

We threshold the singular values to determine the proper dimension. A normalization strategy is used to make the threshold adaptive to the data and noise level. After normalizing all singular values by the mean of 10 larges singular values, the dimension cut is set at where the sum of ten consecutive normalized singular values is smaller than 0.1. We point out that since the spectral embedding is defined by each eigen-function of the LB operator, instead of the collectively spanned space, the accuracy and stability of each singular vectors is more relevant in our case. Perturbation analysis for singular values and vectors in [43] can be used to justify our strategy as well.

(2) **Computing functional map based on anchor pairs.**

Once the proper spectral dimension, k, well spanned by the selected anchor pairs $\{(x_\ell, T(x_\ell) \mid \ell = 1, \cdots, m\}$ is determined, we compute the functional map for k dimensional spectral embedding only based on the anchor pairs. In other words, we consider the following problem

$$\begin{aligned} C &= \arg\min_{C \in \mathfrak{R}} \sum_{\ell=1}^{m} \| \; \Phi^k_{\mathcal{M}_1}(x_\ell) - \Phi^k_{\mathcal{M}_2}(T(x_\ell))C\|^2_F \\ &= \mathrm{Proj}_{\mathfrak{R}} \left(\left(\Phi^k_{\mathcal{M}_1}(\{x_\ell\}^m_{\ell=1})\right)^\top \Phi^k_{\mathcal{M}_2}(\{T(x_\ell)\}^m_{\ell=1}) \right) \end{aligned} \tag{3.5}$$

This restricted version of computing a functional map has twofold advantages. On the one hand, as we discussed in Section 3.1 and appendix, it excludes potential corruption from inaccurate correspondence and leads to better estimation of a functional map. On the other hand, it 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 (3.3), where a KNN search method is applied.

In summary, the full procedure is described in Algorithm 3.1.

Algorithm 3.1 Dual Iterative Refinement (DIR)

```
input Initial correspondence, $K$ LB eigen-functions $\Phi_{\mathcal{M}_1}$ and $\Phi_{\mathcal{M}_2}$ for both shapes, local mapping distortion error threshold $\{\epsilon_i\}_1^N$, singular value threshold, maximum iteration number $N$,
repeat
  1. Find anchor pairs $\{(x_\ell, T(x_\ell) \mid \mathcal{D}_\gamma(T)(x_\ell) < \epsilon_i\}, l = 1, \dots, m_i$ in spatial domain from current correspondence $T$.
  2. Find the proper spectral dimension $k_i$ from the SVD decomposition of the correlation matrix $\Phi^K_{\mathcal{M}_1}(\{x_\ell\}^{m_i}_{\ell=1})^\top \Phi^K_{\mathcal{M}_2}(T(\{x_\ell\}^{m_i}_{\ell=1}))$ of anchor pairs in the spectral domain.
  3. Update functional map $C = \mathrm{Proj}_{\mathfrak{R}} \left( \left(\Phi^{k_i}_{\mathcal{M}_1}(\{x_\ell\}^{m_i}_{\ell=1})\right)^\top \Phi^{k_i}_{\mathcal{M}_2}(\{T(x_\ell)\}^{m_i}_{\ell=1}) \right)$
  4. Update $T(p) = \arg\min_q \|\Phi^{k_i}_{\mathcal{M}_1}(p,:) - \Phi^{k_i}_{\mathcal{M}_2}(q,:)C^\top\|^2$ for all $p \in \mathcal{M}_1$.
until N is reached, or $k_i$ reaches $K$
```

Computation complexity The complexity for checking local mapping distortion is $O(n)$ since the geodesic distance is only computed in a fixed neighborhood, e.g., first ring or second ring, at each point. To compute 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 is used to solve Equation (3.3) which has a complexity of $O(n\log(n))$ [73]. Hence altogether, our method is of complexity $O(n\log(n))$. If we use other global features instead of spectral features, such as geodesic distance to anchor points, it is still $O(n\log(n))$ since we limit the maximal number of anchor points (as the number of LB eigen-functions). The complexity of computing the geodesic distance and solving the resulting assignment problem is still $O(n\log(n))$.

3.4 Discussion

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

3.4.1 Combination of local feature and global feature

For shape correspondence problem, an efficient and robust approach should rely on 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 local mapping distortion criterion is based on local spatial features while the spectral features are global. We would like to point out that the proposed strategy of iterative dual refinement process can be applied to other features as long as they include local and global information complementarily. Other appropriate local and global features can be used in different applications. For examples, one may alternatively use Heat Kernel

Signature [14], Wave Kernel Signature [5] or Geodesic Distance Signature [75] as global features. For shapes with holes and boundaries, we choose Geodesic Distance Signature, which is less sensitive to local mesh distortions and boundaries for most interior points.

3.4.2 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 the morphological study in medical imaging where landmarks could be annotated based on specific tasks [26]. 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 for the algorithm to converge. The final convergence performance of our numerical tests, as shown in Figure 3.7, indicates that DIR can provide accurate correspondence based on only four landmarks, and it is stable with respect to different initialization.

3.4.3 Point cloud matching

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 [37] with adaptive KNN introduced in Algorithm 2.3, to compute the LB eigenvalues and local mapping distortion. DIR works well on matching raw point clouds, as shown in our numerical experiments.

3.4.4 Patch/partial matching

Patch/partial matching often comes with difficulties in artificial boundaries, different sizes, and topological perturbation. Most methods based on LB eigen-functions, such as functional map, kernel matching, heat kernel signature, can not handle these difficulties well. However, our local feature, local mapping distortion at interior points, is not affected by the above difficulties. By replacing spectral features based on LB eigen-functions with geodesic distance to anchor points, which is a global feature less sensitive to the above difficulties for most interior points, DIR can handle patch or partial matching well as shown in our numerical experiments.

3.5 Experiments

In this section, we conduct comprehensive experiments to evaluate the performance of our method on various data sets including TOSCA [12], SCAPE [4], SHREC'16 [19] and SHREC'20 [21] as well as patches extracted from TOSCA. TOSCA and SCAPE are data sets of nearly isometric shapes; SHREC'16 is also a data set of nearly isometric shapes but includes cuts and holes. SHREC'20 is a data set from real-world 3D scans that include various deformation and noise.

In all experiments, no pre-processing, such as a low-resolution model or precomputing 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 that appeared in the papers. Computation using ZoomOut is produced from the code on GitHub shared by the authors [46]. All experiments using our methods are conducted in Matlab with 16GB RAM and Intel i7-6800k CPU.

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

Hyperparameters We set neighborhood size for local mapping distortion criterion as the second ring; maximum iteration number is 10; and the local mapping distortion threshold is set to [0.26, 0.22, 0.18, 0.14, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1] at each iteration. The initial guess is given by the KNN search based on SHOT features [71]. These hyperparameters, especially the local mapping distortion, are not sensitive to different data sets. We use the same hyperparameters on TOSCA and SCAPE data sets. For SHREC'16, SHREC'20 and patch matching, we set the maximum number of anchor pairs served as the reference points for geodesic distance feature to 800.

TOSCA and SCAPE TOSCA data set consists of 76 shapes in 8 different categories. The vertex number is from 4k to 50k. It contains both human and animal shapes. We present the results of three experiments on TOSCA data. The first one uses point clouds with its original mesh as input, and computes 1000 spectral basis as the maximal spectral dimension; the second one also uses point clouds with its original mesh as input but computes 500 spectral basis as the maximal spectral dimension; the last one uses only point clouds (no mesh) as input, and computes 1000 spectral basis (using local mesh method) as the maximal spectral dimension. SCAPE data set has 72 shapes of the same person with different poses, and each contains 12,500 vertices. We perform the same experiments as on TOSCA data.

The results on TOSCA and SCAPE data sets are illustrated in Figure 3.3. We compare our method with the following methods: Blended [36], Best Conformal [36], GMDS [11], Kernel Marching [76], SEQA [80] and RSWD [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 high accuracy.

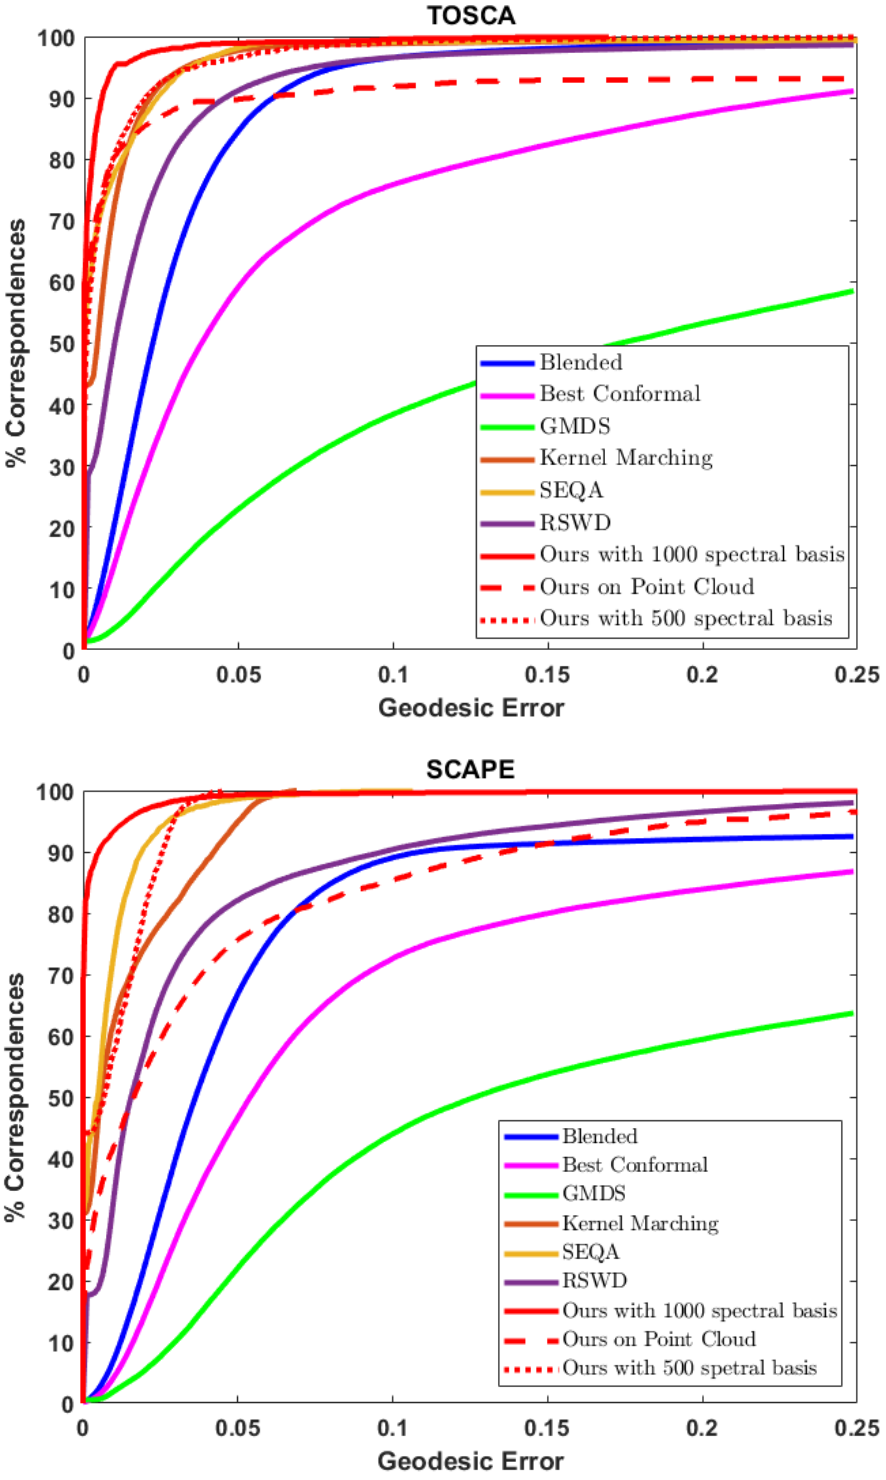

Figure 3.3: Geodesic error on the TOSCA and SCAPE data sets with comparison to other methods.

SHREC'16 SHREC'16 Partial Correspondence 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.

The results are illustrated in Figure 1.4. We compare our method with ZoomOut [46], Partial Functional Maps [61] and Random Forests [62]. Since SHREC'16 only provides ground truth map with barycentric coordinates, we compute the geodesic error at each point as the weighted average of geodesic errors to the three vertices in the target triangle element. The results show that our method is quite flexible and robust to handle shape matching for different scenarios. It again outperforms other state-of-the-art methods markedly and achieves high accuracy on this challenging data set.

SHREC'20 SHREC'20 is a data set from 3D scans of a real-world stuffed toy rabbit that contains deformations caused by stretching, indent, twisting and inflation. To make it even more challenging, the rabbit is inflated by different internal materials, and there are partial scans and missing parts caused by occlusion. Since SHREC'20 does not provide ground truth, hence we only show the visual result as illustrated in Figure 3.4. We also present final local mapping distortion heat maps to indicate local deformation at each point. These heat maps indicate that our method can still produce accurate results, although the target shape may have nontrivial metric deformation from the source shape. Moreover, the local mapping distortion provides an excellent measurement to quantify the mapping quality without knowing the ground truth correspondence. As we can observe from the second row in Figure 3.4, the mapping quality of shapes in the second and fourth columns is satisfactory, while the third one has relatively large distortion in some regions. This is because the target shape is deformed from the source shape by a very distorted metric deformation. In our future work, we would extend our investigation to find high-quality mappings between

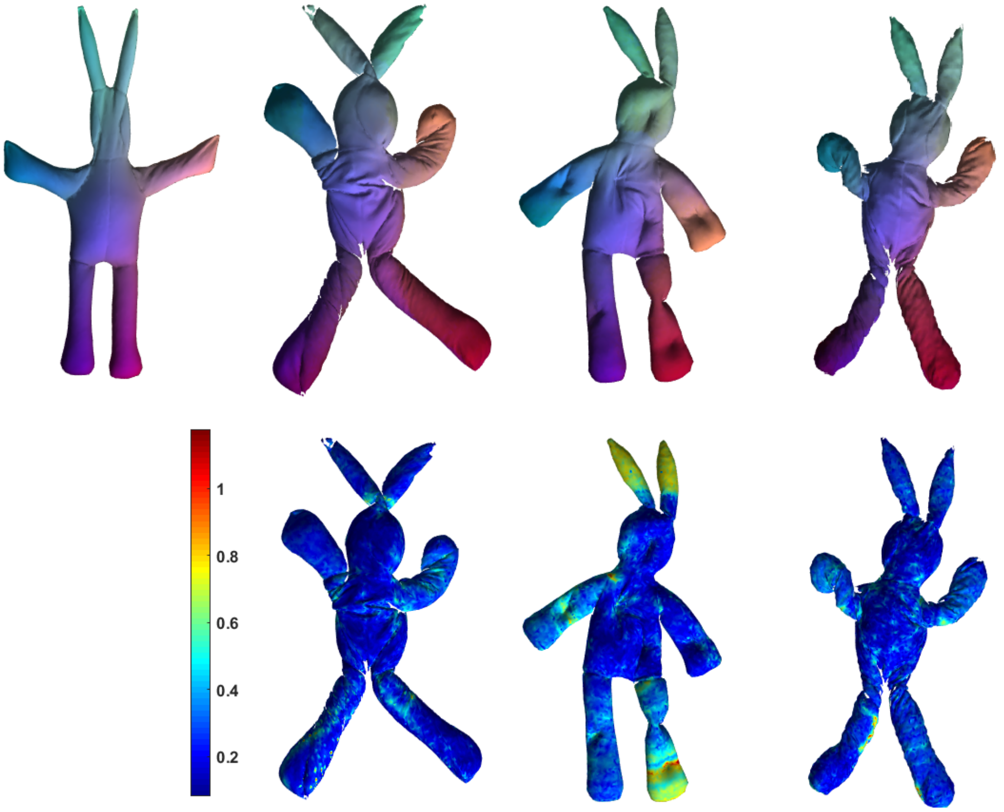

Figure 3.4: Qualitative result of our method on SHREC'20. In the first row, the rabbit in the first column is mapped to the rest rabbits in the first row. In the second row, final local mapping distortion is computed on the target shapes as a heat map. It is clear that the area where local mapping distortion is lower shows a better visual mapping result.

shapes with large intrinsic deformations.

Patch Matching Since there is no standard data set for patch matching, so instead, we report several experiments using patches cut from TOSCA data set. The first case contains two examples with artificial boundaries, different sizes (partial patching), and topological changes simultaneously. In the first example, we take a portion of an arm (fingertips also removed) from one of two 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 mesh surface with no holes, while the arm and the body patch are not closed and have holes and boundaries. Our simple method again performs well for this challenging example, as shown in Figure 3.5.

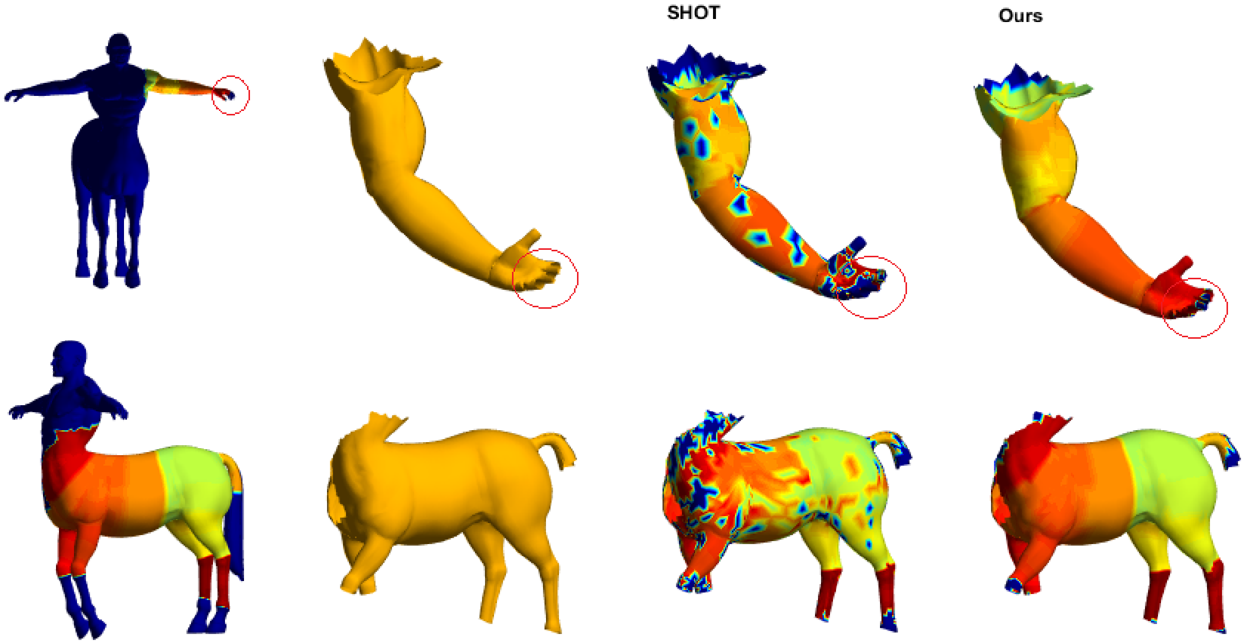

Figure 3.5: Examples of patch matching. Removed fingertips are highlighted by red circles. The first column is the entire centaur, and its non-blue area is the ground truth color map. Extra points are colored in blue. The arm or body part in the second column is mapped onto the entire shape. The third column is the color mapping result using SHOT, and the last column is the color mapping result using our method.

The second case is matching two patches that contain 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 local mapping distortion test is added to prune out those points. The result is shown in Figure 3.6. Our method still performs quite well. These experiments on patch matching indicate that our method is robust to size differences, artificial boundaries, and topological changes.

Experiments Using Limited Landmarks We select 4 pairs of centaur shapes from TOSCA, initialize our method with a different number of randomly chosen landmarks (without using SHOT for initialization) and then plot the average geodesic error curve. After a different number of iterations, the final performance is illustrated in Figure 3.7. Even with four initial random landmarks, the performance is super, although more iterations are

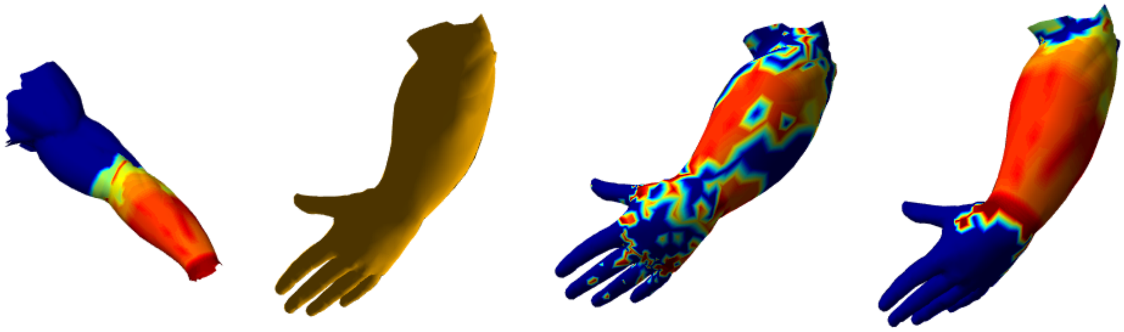

Figure 3.6: Examples of patch matching. Patch in the second column is mapped to the patch in the first column. Extra points in the first patch are colored in blue, and its non-blue area is the ground truth color map. The third column shows the mapping result using SHOT, and the last column is the mapping result using our method. The blue color in column three and four indicates points which do not pass the post local mapping distortion test.

needed. It shows the stability of our method.

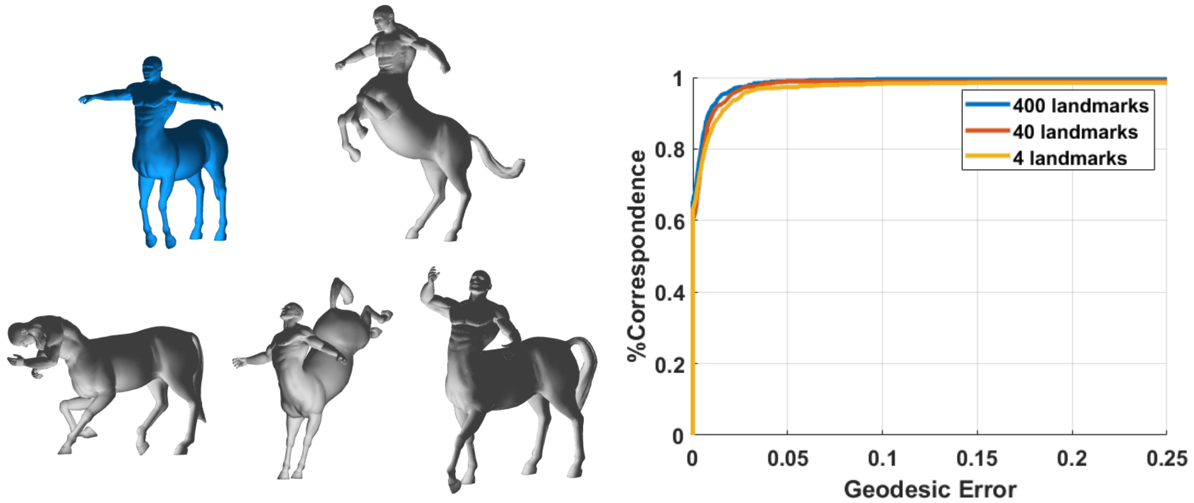

Figure 3.7: Average convergence performance of our method testing on several pairs of centaur shapes from TOSCA (blue centaur is mapped to gray centaurs), given different number of initial random annotated landmarks.

Run Time In Table 3.1, we show the average run time of our method on several examples from TOSCA data set, which includes shapes with vertices ranging from 4,000 to more than 50,000 and indicates the computation efficiency of our methods. As mentioned in Section 3.3, the complexity of our method is $O(n\log(n))$. However, in practice, when computing point clouds 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

Model	Wolf	Centaur	Horse	Cat	David
Number of Vertices	4344	15768	19248	27894	52565
SEQA	59	531	801	929	1681
Kernel Matching	60	NA	NA	NA	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 3.1: Average run time for shapes from TOSCA

David example 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 that report the time.

3.6 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 (selected from local mapping distortion using local features) to guide the next refinement (using spectral or other global features), 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-art methods in terms of accuracy, efficiency, and stability.

Bibliography

[1] Y. Aflalo, A. Bronstein, and R. Kimmel. Graph matching: relax or not? *arXiv preprint arXiv:1401.7623*, 2014.

[2] Y. Aflalo, A. Bronstein, and R. Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015.

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

[4] D. Anguelov, P. Srinivasan, H.-C. Pang, D. Koller, S. Thrun, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in neural information processing systems*, pages 33–40, 2005.

[5] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 1626–1633. IEEE, 2011.

[6] J. Barzilai and J. M. Borwein. Two-point step size gradient methods. *IMA journal of numerical analysis*, 8(1):141–148, 1988.

[7] P. Bérard, G. Besson, and S. Gallot. Embedding riemannian manifolds by their heat kernel. *Geometric & Functional Analysis GAFA*, 4(4):373–398, 1994.

[8] D. P. Bertsekas. *Network optimization: continuous and discrete models*. Athena Scientific Belmont, 1998.

[9] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016.

[10] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal on Scientific Computing*, 28(5):1812–1836, 2006.

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

[12] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes.* Springer Science & Business Media, 2008.

[13] 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.

[14] M. M. Bronstein and I. 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.

[15] J. Cates, P. T. Fletcher, M. Styner, M. Shenton, and R. Whitaker. Shape modeling and analysis with entropy-based particle systems. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 333–345. Springer, 2007.

[16] I. Chavel. *Eigenvalues in Riemannian geometry*, volume 115. Academic press, 1984.

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

[18] D. Conte, P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 18(03):265–298, 2004.

[19] L. Cosmo, E. Rodolà, J. Masci, A. Torsello, and M. Bronstein. Matching deformable objects in clutter. In *Proc. 3D Vision (3DV)*, pages 1–10, 2016.

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

[21] R. Dyke. *Non-Rigid Shape Correspondence of Physically Based Deformations*, 2020.

[22] N. Dym, H. Maron, and Y. Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv preprint arXiv:1705.06148*, 2017.

[23] R. A. Earnshaw. *Virtual reality systems.* Academic press, 2014.

[24] A. Gasparetto, L. Cosmo, E. Rodola, M. Bronstein, and A. 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.

[25] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Transactions on pattern analysis and machine intelligence*, 18(4):377–388, 1996.

[26] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE transactions on medical imaging*, 23(8):949–958, 2004.

[27] X. D. Gu, R. Guo, F. Luo, and W. Zeng. Discrete laplace-beltrami operator determines discrete riemannian metric. *arXiv preprint arXiv:1010.4070*, 2010.

[28] S. Gumhold, X. Wang, and R. S. MacLeod. Feature extraction from point clouds. In *IMR*. Citeseer, 2001.

[29] O. Halimi, O. Litany, E. Rodolà, A. Bronstein, and R. Kimmel. Self-supervised learning of dense shape correspondence. *arXiv preprint arXiv:1812.02415*, 2018.

[30] L. Han, S. Liu, B. Yu, S. Xu, and R. Xiang. Orientation-preserving spectral correspondence for 3d shape analysis. *Journal of Imaging Science and Technology*, 2019.

[31] V. Jain and H. 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.

[32] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal of Shape Modeling*, 13(01):101–124, 2007.

[33] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.

[34] I. Kezurer, S. Z. Kovalsky, R. Basri, and Y. Lipman. Tight relaxation of quadratic matching. In *Computer Graphics Forum*, volume 34, pages 115–128. Wiley Online Library, 2015.

[35] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. In *ACM Transactions on Graphics (TOG)*, volume 26, page 64. ACM, 2007.

[36] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, page 79. ACM, 2011.

[37] R. Lai, J. Liang, and H. Zhao. A local mesh method for solving pdes on point clouds. *Inverse Problems & Imaging*, 7(3), 2013.

[38] R. Lai, Y. Shi, K. Scheibel, S. Fears, R. Woods, A. W. Toga, and T. 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.

[39] R. Lai and H. 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.

[40] E. L. Lawler. The quadratic assignment problem. *Management science*, 9(4):586–599, 1963.

[41] M. Leordeanu and M. 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.

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

[43] J. Liu, X. Liu, and X. Ma. First-order perturbation analysis of singular vectors in singular value decomposition. *IEEE Transactions on Signal Processing*, 56(7):3044–3049, 2008.

[44] S. Marschner and P. Shirley. *Fundamentals of computer graphics.* CRC Press, 2015.

[45] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. 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.

[46] S. Melzi, J. Ren, E. Rodolà, A. Sharma, P. Wonka, and M. Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):155, 2019.

[47] F. Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11(4):417–487, 2011.

[48] M. E. Mortenson. Geometric modeling. 1997.

[49] M. Nixon and A. Aguado. *Feature extraction and image processing for computer vision.* Academic Press, 2019.

[50] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012.

[51] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. In *Computer Graphics Forum*, volume 29, pages 1555–1564. Wiley Online Library, 2010.

[52] M. Ovsjanikov, Q. Merigot, F. Memoli, and L. Guibas. One point isometric matching with the heat kernel. *CGF*, 29(5):1555–1564, 2010.

[53] J. R. Parker. *Algorithms for image processing and computer vision.* John Wiley & Sons, 2010.

[54] A. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modeling and recognition. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (7):715–729, 1991.

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

[56] J. N. Reddy. An introduction to the finite element method. *New York*, 1993.

[57] F. Rengier, A. Mehndiratta, H. Von Tengg-Kobligk, C. M. Zechmann, R. Unterhinninghofen, H.-U. Kauczor, and F. L. Giesel. 3d printing based on imaging data: review of medical applications. *International journal of computer assisted radiology and surgery*, 5(4):335–341, 2010.

[58] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete laplace–beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381–390, 2009.

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

[60] E. Rodola, A. M. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 182–189. IEEE, 2012.

[61] E. Rodolà, L. Cosmo, M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Computer Graphics Forum*, 36(1):222–236, 2017.

[62] E. Rodolà, S. Rota Bulo, T. Windheuser, M. Vestner, and D. Cremers. Dense non-rigid shape correspondence using random forests. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4184, 2014.

[63] Y. Sahillioğlu. A genetic isometric shape correspondence algorithm with adaptive sampling. *ACM Transactions on Graphics (TOG)*, 37(5):1–14, 2018.

[64] S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM (JACM)*, 23(3):555–565, 1976.

[65] S. C. Schonsheck, M. M. Bronstein, and R. Lai. Nonisometric surface registration via conformal laplace-beltrami basis pursuit. *arXiv preprint arXiv:1809.07399*, 2018.

[66] S. Skiena. Involutions. *Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*, pages 32–33, 1990.

[67] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, volume 28, pages 1383–1392. Wiley Online Library, 2009.

[68] R. Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.

[69] A. A. Taha and A. Hanbury. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC medical imaging*, 15(1):29, 2015.

[70] A. Tevs, A. Berner, M. Wand, I. 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.

[71] F. Tombari, S. Salti, and L. Di Stefano. Unique signatures of histograms for local surface description. In *European conference on computer vision*, pages 356–369. Springer, 2010.

[72] L. N. Trefethen and D. Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.

[73] P. M. Vaidya. An o (n logn) algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4(2):101–115, 1989.

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

[75] O. Van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011.

[76] M. Vestner, Z. Lähner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodola, A. Bronstein, M. Bronstein, R. Kimmel, et al. Efficient deformable shape correspondence via kernel matching. In *3D Vision (3DV), 2017 International Conference on*, pages 517–526. IEEE, 2017.

[77] M. Vestner, R. Litman, E. Rodolà, A. Bronstein, and D. Cremers. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3327–3336, 2017.

[78] C. Wang, M. M. Bronstein, A. M. Bronstein, and N. 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.

[79] A. Watt. 3d computer graphics. 1993.

[80] R. Xiang, R. Lai, and H. Zhao. Efficient and robust shape correspondence via sparsity-enforced quadratic assignment. *arXiv preprint arXiv:2003.08680*, 2020.

[81] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. Van Kaick, and A. Tagliasacchi. Deformation-driven shape correspondence. *Computer Graphics Forum*, 27(5):1431–1439, 2008.