

Symmetry-Aware Nonrigid Matching of Incomplete 3D Surfaces

Yusuke Yoshiyasu¹, Eiichi Yoshida^{1,2}, Kazuhito Yokoi¹, Ryusuke Sagawa¹

¹ Intelligent Systems Research Institute ² CNRS-AIST JRL (Joint Robotics Laboratory), UMI3218/CRT
National Institute of Advanced Industrial Science and Technology (AIST), Japan
{yusuke-yoshiyasu, e.yoshida, Kazuhito.Yokoi, ryusuke.sagawa}@aist.go.jp

Abstract

We present a nonrigid shape matching technique for establishing correspondences of incomplete 3D surfaces that exhibit intrinsic reflectional symmetry. The key for solving the symmetry ambiguity problem is to use a point-wise local mesh descriptor that has orientation and is thus sensitive to local reflectional symmetry, e.g. discriminating the left hand and the right hand. We devise a way to compute the descriptor orientation by taking the gradients of a scalar field called the average diffusion distance (ADD). Because ADD is smoothly defined on a surface, invariant under isometry/scale and robust to topological errors, the robustness of the descriptor to non-rigid deformations is improved. In addition, we propose a graph matching algorithm called iterative spectral relaxation which combines spectral embedding and spectral graph matching. This formulation allows us to define pairwise constraints in a scale-invariant manner from k -nearest neighbor local pairs such that non-isometric deformations can be robustly handled. Experimental results show that our method can match challenging surfaces with global intrinsic symmetry, data incompleteness and non-isometric deformations.

1. Introduction

Finding meaningful correspondences between two shapes is an important task in computer vision and computer graphics. If the correspondences between surfaces are known, we can align 3D scans [12], morph two shapes, establish statistical models [2] and transfer various types of information—such as textures, segmentations and even deformations [30]—from one surface to another.

Early approaches focused on a rigid case i.e., establishing correspondences between two nearly isomorphic objects (two objects with almost the same shape). In this case, because a mapping between two shapes can be parameterized by few parameters (rotation and translation), optimization is achieved efficiently by the iterative closest point (ICP) technique and graph matching algorithms [12, 33].

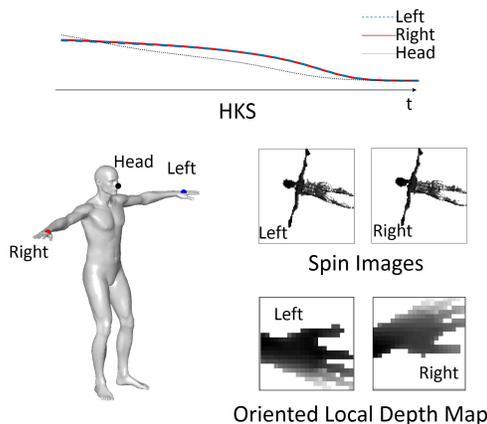


Figure 1. Comparisons of local descriptors. Although heat kernel signatures (HKS) are multi-scale, they do not discriminate reflectional symmetry pairs, e.g., the left hand and the right hand. Spin images are also not discriminative. In contrast, the oriented local depth maps can differentiate.

Establishing correspondences of two shapes undergoing nonrigid deformations, however, can no longer be represented by simple parameters. Consequently, most of the methods represent a correspondence using a point-to-point (or part-to-part) match and must solve a large combinatorial optimization problem. Existing techniques [5, 19, 25, 29, 37] try to find correspondences by minimizing some structural distortions that are defined as a point-wise cost (differences of feature descriptors) [3, 8, 32], a pairwise cost (differences of distances) [3, 13] and a high-order cost (triplets) [36].

Although many nonrigid matching techniques have been proposed, this problem is still very difficult due to the following challenges: The first challenge is *the symmetric ambiguity problem*, i.e., the ambiguity that arises when matching the surfaces with intrinsic reflectional symmetry. The second challenge is *incompleteness* of input surfaces. Real-world 3D scans are incomplete and have holes/noises. The third challenge is *non-isometric deformations*. Isometric assumptions are often used in nonrigid shape matching but it

is in general not the case that two surfaces that we would like to match are isometric.

Previous work solves the symmetry ambiguity problem by incorporating the distortion measures based on conformal mapping [17, 19], high-order potentials [36] or symmetry information [20, 24]. However, these solutions are only applicable to limited classes of surfaces i.e., genus-zero closed surfaces (or, at least, holes of two surfaces must be located at same places), which cannot be used for matching real-world 3D scans that are incomplete. Several techniques have been proposed to tackle the incompleteness problem [8, 13, 32], which is typically accomplished by constructing a pairwise cost from the distance that is robust to topological noises. However, this approach is only limited to match surfaces that have undergone isometric deformation. The point-wise descriptors that have been incorporated in nonrigid matching so far—such as spin images [15], curvatures or heat kernel signatures [31]—is robust to data incompleteness. Some of them are also invariant under isometry and local scales [6]. However, they do not have orientations such that they cannot evaluate local symmetries and are not effective in discriminating global symmetry pairs.

In this paper, we propose a nonrigid matching algorithm for establishing correspondences of incomplete 3D surfaces in the presence of intrinsic symmetries. We overcome 1) symmetry ambiguity, 2) data incompleteness and 3) non-isometric deformations at the same time, which is not yet achieved by others [17, 26, 32, 36] but is actually very important for non-rigid 3D scan registration, e.g., automatically matching a template mesh to 3D scans with different details, local scales and poses. The contributions of this paper are the following:

- We propose the idea of using orientation-aware local mesh descriptors [9, 22, 35] to match two surfaces with intrinsic symmetry. To our knowledge, this is the first work that uses this class of point-wise descriptors to solve symmetry ambiguity during nonrigid shape matching. Unlike triplet descriptors [17, 36], this strategy is more efficient and robust to data incompleteness.
- We propose a new way of computing descriptor orientation by taking gradients of the average diffusion distance (ADD) field. Since ADD is smooth, invariant under scale/isometry and robust to topological errors, the descriptor orientation can be computed consistently for two different surfaces.
- We propose a graph matching technique called iterative spectral relaxation that combines spectral embedding and spectral graph matching. The benefit of our algorithm is its ability to attain global convergence with the pairwise cost that evaluates local pairwise matches only for K -nearest neighbors. Thus, it is

efficient. More important, unlike [13, 26, 32] which evaluates geodesic distortions, this formulation makes the pairwise cost robust to sampling inconsistency and non-isometric deformations.

Note that, to handle data incompleteness, we use diffusion distances for defining point-wise/pairwise constraints.

2. Background

Non-rigid surface matching can be formulated as a graph matching problem that incorporates point-wise costs and pairwise costs [11], which is a quadratic assignment problem (QAP). Here, we briefly review basics of graph matching, QAP and previous relaxation techniques that are frequently used in 3D shape matching. Please refer to e.g., [7, 38] for detailed surveys.

2.1. Graph matching and QAP

The goal of graph matching is to establish point-to-point correspondences between point $i = 1, 2, \dots, n$ on shape \mathcal{S}_1 and point $k = 1, 2, \dots, m$ on shape \mathcal{S}_2 . Point-to-point correspondences can be expressed with a permutation matrix, $\mathbf{X} \in \{0, 1\}^{n \times m}$ [7]:

$$\mathbf{X}(i, k) = \begin{cases} 1 & i \text{ corresponds to } k \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$\text{s.t. } \mathbf{X}\mathbf{1}_m \leq \mathbf{1}_n, \mathbf{X}^T\mathbf{1}_n = \mathbf{1}_m \quad (2)$$

The inequality in the above constraint is used for the case when \mathcal{S}_1 and \mathcal{S}_2 are of different sizes. We assume $n > m$ throughout this paper.

Graph matching seeks for correspondences that minimize some structural distortions or, conversely, that maximize some structural affinity measures. Let \mathbf{f}_i and \mathbf{g}_k be N dimensional vectors that contain feature descriptor values of point i and point k . Then, the point-wise affinity measure is defined as:

$$K_p(i, k) = \exp(-\|\mathbf{f}_i - \mathbf{g}_k\|/\sigma_p) \quad (3)$$

Next, let $d(i, j)$ and $d(k, l)$ be distances between points (i, j) on \mathcal{S}_1 and points (k, l) on \mathcal{S}_2 . Then, the pairwise affinity measure is defined as:

$$K_q(i, j, k, l) = \exp(-\|d(i, j) - d(k, l)\|/\sigma_q) \quad (4)$$

Here, σ_p and σ_q are scale factors.

Now that point-wise and pairwise affinities, K_p and K_q , are defined, the graph matching problem is formulated as QAP:

$$\begin{aligned} \operatorname{argmax}_{\mathbf{X}} & \sum_{i,k} K_p(i, k)\mathbf{X}(i, k) \\ & + \sum_{i,j,k,l} K_q(i, j, k, l)\mathbf{X}(i, k)\mathbf{X}(j, l) \end{aligned} \quad (5)$$

In the graph matching field, two major formulations for QAP have been introduced: Lawler’s formulation and Koopmans-Beckmann’s formulation.

Lawler’s QAP Let us define a global affinity matrix $\mathbf{K} \in \mathbb{R}^{nm \times nm}$ that contains point-wise affinity $K_p(i, k)$ at its diagonal and the pairwise affinity $K_q(i, j, k, l)$ at off-diagonals, where its row a and column b correspond to i, j, k and l as $a = n(k - 1) + i$ and $b = n(l - 1) + j$. Then, the entry of \mathbf{K} is:

$$\mathbf{K}(a, b) = \begin{cases} K_p(i, k) & b = a \\ K_q(i, j, k, l) & \text{otherwise} \end{cases} \quad (6)$$

With \mathbf{K} , the problem Eq. 5 can be written as:

$$\operatorname{argmax}_{\mathbf{X}} \operatorname{vec}(\mathbf{X})^T \mathbf{K} \operatorname{vec}(\mathbf{X}) \quad (7)$$

where $\operatorname{vec}(\mathbf{X})$ is a column vector of \mathbf{X} , $\operatorname{vec}(\mathbf{X}) = [x_{11}, x_{21}, \dots, x_{n1}, x_{12}, \dots, x_{nm}]^T$.

Koopmans-Beckmann’s QAP Let $\mathbf{F} \in \mathbb{R}^{n \times n}$ and $\mathbf{G} \in \mathbb{R}^{m \times m}$ be adjacency matrices whose entries are $F_{ij} = d(i, j)$ and $G_{kl} = d(k, l)$, respectively. The problem Eq. 5 can alternatively be written as [7, 33]:

$$\operatorname{argmin}_{\mathbf{X}} \|\mathbf{X}\mathbf{F}\mathbf{X}^T - \mathbf{G}\|^2 + \omega \|\mathbf{X}\mathbf{f} - \mathbf{g}\|^2 \quad (8)$$

where \mathbf{f} and \mathbf{g} are $n \times N$ and $m \times N$ matrices containing N -D feature descriptor values. In addition, ω balances the constraints.

2.2. Relaxation techniques

The problem Eq. 5 is NP-hard and is highly expensive to solve with an exact optimal algorithms like branch-and-bound. Thus, many approximation techniques have been proposed. Among them, we focus on spectral relaxation techniques that are commonly used for 3D shape matching: spectral graph matching [18] and spectral embedding [33].

Spectral graph matching [18] This algorithm is a relaxed version of Lawler’s QAP (Eq. 7). It relaxes the constraint on permutation matrix \mathbf{X} to $\|\operatorname{vec}(\mathbf{X})\|^2 = 1$, which results in maximization of Rayleigh quotient:

$$\operatorname{argmax}_{\mathbf{X}} \frac{\operatorname{vec}(\mathbf{X})^T \mathbf{K} \operatorname{vec}(\mathbf{X})}{\operatorname{vec}(\mathbf{X})^T \operatorname{vec}(\mathbf{X})} \quad (9)$$

Thus, \mathbf{X} is given by the eigen vector corresponding to the largest eigen value of \mathbf{K} , which defines scores for correspondences. The eigen vector is then binarized with the Hungarian algorithm.

The advantage of this approach is its global convergence. In the context of nonrigid shape matching in 3D, Huang et

al. [13] and Kim et al. [17] used this strategy to compute confidence scores of matching. The difficulty of using this method is memory requirements and computational costs for construction and eigen decomposition of \mathbf{K} . The complexity for constructing \mathbf{K} is $O(n^2, n^2)$; for instance, when matching 100 points to 100 points, it will reach 10^8 .

Spectral embedding [33] This algorithm is a relaxed version of Koopmans-Beckmann’s QAP (Eq. 8). It relaxes constraints on permutation matrix \mathbf{X} , such that \mathbf{X} is approximated as an orthogonal matrix, $\mathbf{X}^T \mathbf{X} = \mathbf{X} \mathbf{X}^T = \mathbf{I}$. Under this assumption, Eq. 8 is reduced to a linear assignment problem that establishes correspondences between the eigen vectors of \mathbf{F} and \mathbf{G} . However, due to numerical approximations, sign flips and ordering switches of the eigen vectors occur [14, 21]. Let \mathbf{U} and \mathbf{V} be the eigen vectors of \mathbf{F} and \mathbf{G} ¹. Let \mathbf{S} and \mathbf{O} be matrices containing signs and ordering of eigen vectors. Then, we can rewrite Eq. 8 as:

$$\operatorname{argmin}_{\mathbf{X}, \mathbf{S}, \mathbf{O}} \|\mathbf{X}\mathbf{V} - \mathbf{U}\mathbf{S}\mathbf{O}\|^2 + \omega \|\mathbf{X}\mathbf{f} - \mathbf{g}\|^2 \quad (10)$$

Further, if we express \mathbf{S} and \mathbf{O} using a single orthonormal matrix, $\mathbf{R} = \mathbf{S}\mathbf{O}$, this leads to:

$$\operatorname{argmin}_{\mathbf{X}, \mathbf{R}} \|\mathbf{X}\mathbf{V} - \mathbf{U}\mathbf{R}\|^2 + \omega \|\mathbf{X}\mathbf{f} - \mathbf{g}\|^2 \quad (11)$$

Equation 11 is solved via an ICP-like alternating optimization approach [21, 23, 26]. Given an initial guess of correspondences, orthonormal matrix \mathbf{R} and permutation matrix \mathbf{X} are iteratively obtained until convergence. The first step obtains \mathbf{R} with \mathbf{X} fixed:

$$\operatorname{argmin}_{\mathbf{R}} \|\mathbf{V}' - \mathbf{U}\mathbf{R}\|^2 \text{ s.t. } \mathbf{R}^T \mathbf{R} = \mathbf{R}\mathbf{R}^T = \mathbf{I} \quad (12)$$

where $\mathbf{V}' = \mathbf{X}^0 \mathbf{V}$ with \mathbf{X}^0 indicating the permutation matrix computed in the previous step. To solve Eq. 12, \mathbf{R} is approximated as $\mathbf{R}^* = \mathbf{V}' \mathbf{U}^T$ and then ortho-normalized with SVD. The second step obtains \mathbf{X} with \mathbf{R} fixed:

$$\operatorname{argmin}_{\mathbf{X}} \|\mathbf{X}\mathbf{V} - \mathbf{U}'\|^2 + \omega \|\mathbf{X}\mathbf{f} - \mathbf{g}\|^2 \quad (13)$$

where $\mathbf{U}' = \mathbf{U}\mathbf{R}^0$ with \mathbf{R}^0 indicating the orthonormal matrix computed in the previous step. This is actually a linear assignment problem. Thus, \mathbf{X} is computed with the Hungarian algorithm.

The advantage of this technique is that it is efficient and requires a small memory space. The drawback of this approach is that it is local, i.e., it requires a good initial guess or otherwise it will be trapped by local minimum. Furthermore, ortho-normalization of \mathbf{R} is too weak to achieve

¹Note that recent techniques [21, 23, 26] perform eigen decomposition on the Laplace-Bertrami operator, instead of on a dense adjacency matrix. This is due to the fact that the Laplace-Bertrami operator is sparse and can be efficiently computed.

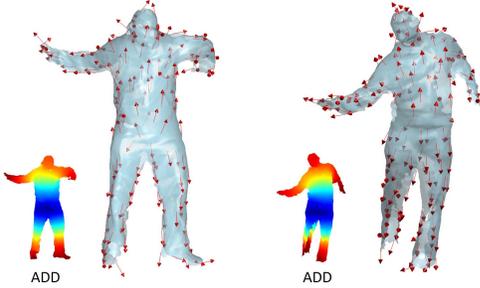


Figure 2. Gradient vectors of average diffusion distances (ADD) are isometric invariant and smooth.

isometry because $\mathbf{R} = \mathbf{S}\mathbf{O}$ is usually nearly diagonal when deformation is close to isometry. For this reason, the sparse modeling approach [26] further constrains \mathbf{R} to be approximately diagonal. However, this approach limits the robustness to non-isometric deformations.

3. Symmetry-aware surface matching

Our algorithm takes two triangle meshes \mathcal{M}_1 and \mathcal{M}_2 as inputs, where \mathcal{M}_2 is a surface that have undergone non-rigid deformation from \mathcal{M}_1 . \mathcal{M}_1 and \mathcal{M}_2 consist of $n_{\mathcal{M}_1}$ and $n_{\mathcal{M}_2}$ vertices, respectively. We assume that the mesh is a single component having no isolated pieces, but it can be incomplete and can contain holes. Also, \mathcal{M}_1 and \mathcal{M}_2 may exhibit intrinsic symmetry.

The goal is to match these challenging input surfaces semantically correctly and produce dense correspondences between them. Using the notations in Section 2.1, our tasks are thus to define \mathbf{f} , \mathbf{g} , \mathbf{F} and \mathbf{G} in a proper way and to devise an efficient matching algorithm to solve Eq. 5.

Our algorithm starts from sampling to reduce possible matches (Section 3.1). We employ the farthest point sampling strategy based on the diffusion distance that is robust to topological noises. We next compute an oriented local descriptor at sample points (Section 3.2). This is the key to solving the symmetric ambiguity problem. Finally, we match the points based on an efficient graph matching algorithm (Section 3.3).

3.1. Sampling

To reduce the number of possible matches, we select several hundred (typically 200–300) of points from the original mesh vertices. Following Raviv et al. [27] who achieved symmetry detection of surfaces with topological noises, we employ the farthest point sampling technique based on the diffusion distance [10] to accommodate incomplete surfaces. As opposed to the geodesic distance, the diffusion distance is robust to noise and topological short-circuits.

3.2. Feature description

Once we have obtained sample points, we compute a local descriptor at every point. We propose to discriminate global symmetry pairs based on a the orientation-aware mesh descriptor that is sensitive to local symmetry.

The descriptors that are used in previous nonrigid shape matching techniques [3, 8, 13, 23, 32], e.g., spin images [15], Gaussian curvature and heat kernel signature [31], do not have orientations. Thus, these descriptors are not sensitive to local symmetry e.g., the left hand and the right hand as depicted in Fig. 1. In contrast, recent techniques [9, 22, 35] compute a dominant angle or principal axis from local geometry and align histogram bins to it. Our key observation is that these oriented descriptors can distinguish global reflectional symmetry pairs based on local symmetric information as depicted in Fig. 1.

Robust descriptor orientation In this paper, we extend the descriptor based on the local depth map [9, 22]. The problem of these descriptors when applied to nonrigid shape matching is that the computations of the principal axis using PCA become unstable, due to local deformations. To alleviate this problem, we compute the orientation by taking gradients of a global scalar field that is isometric-invariant and that is stable under local geometric changes. For this, we use an average diffusion distance field (ADD) [10]. ADD is a scalar field whose values are low at the center of the segment and high at the extrema points (Figure 2). Since the diffusion distance is isometric invariant, so is ADD. By using a sufficient high value for the time factor of the diffusion distance, we can make ADD smooth and robust to local geometric changes.

Once we have computed the orientation, we obtain local depth maps. At each sample point, we compute the local depth map of 50×50 pixels as follows:

the view point is located at the position distant from point i by some distance α in the direction of the surface normal; the view direction is the opposite direction of the surface normal; the view-up direction is the normalized gradient vector of ADD. To construct a descriptor, we convert each local depth map to a column vector. To achieve scale-invariance, we normalize the sizes of the input surfaces by the scale factor that normalizes one of the surfaces' bounding-box diagonal to 2. In addition, we achieve multi-scale descriptions by computing local depth maps at three different α values and concatenate three depth maps into one vector.

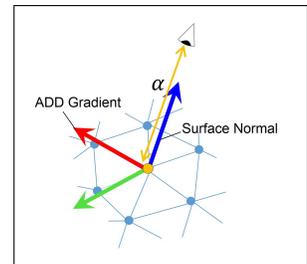




Figure 3. The pairwise affinity is computed from the edges emanating from a vertex and its K nearest neighbor points. By exploiting the sparsity of the matrix we can reduce memory requirements.

3.3. Matching

We feed local descriptors into a graph matching framework that incorporates pairwise constraints. To achieve this efficiently, we propose an algorithm called iterative spectral relaxation that fuses spectral graph matching [18] and spectral embedding [33]. Due to the ability of this algorithm to preserve global consistency of the graph, we can make the pairwise affinity matrix sparse such that we can achieve global convergence while maintaining efficiency.

Iterative spectral relaxation (ISR) Instead of alternating Eq. 12 and Eq. 13 like the spectral embedding approach, we alternate Eq. 12 and Eq. 9. We solve Eq. 9 with \mathbf{U} and \mathbf{V} plugged into the diagonal of \mathbf{K} :

$$\mathbf{K}(a, a) = \exp(-\|\mathbf{f}_i - \mathbf{g}_k\|/\sigma_p) + \exp(-\|\mathbf{V}_i - \mathbf{U}'_k\|/\sigma_e) \quad (14)$$

where $a = n(k-1) + i$. Here, we compute \mathbf{U} and \mathbf{V} from the Laplace-Bertrami operator. Let Φ and Ψ be $M \times n_{\mathcal{M}_1}$ and $M \times n_{\mathcal{M}_2}$ matrices containing the first M dimensional Laplace eigen basis of \mathcal{M}_1 and \mathcal{M}_2 . Then, \mathbf{U} and \mathbf{V} are defined as $\mathbf{U}_i = \Phi_{\text{idx}(i)}$ and $\mathbf{V}_k = \Psi_{\text{idx}(k)}$ where $\text{idx}(i)$ and $\text{idx}(k)$ is indices of sample points.

The advantage of ISR is that it can achieve (near-) global convergence while reducing memory requirements. This is due to the latter term of Eq. 14 that maintains global consistency of the graph. Thus, we only need to evaluate affinities for local edges near sample points and therefore we can make \mathbf{K} sparse. To find point pairs for establishing local edges, we search K nearest neighbor points around point i in terms of the diffusion distance (Fig. 3). We then construct edges that emanate from i to K nearest neighbor points. If edge i - j and edge l - k are in the local edge set, we set $\mathbf{K}(a, b)$ to 1, otherwise 0:

$$\mathbf{K}(a, b) = \begin{cases} 1 & k \in \text{KNN}(i) \cap l \in \text{KNN}(j) \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

where $\text{KNN}(i)$ is a set of K nearest points of point i . This rather rough binary representation allows us to robustly

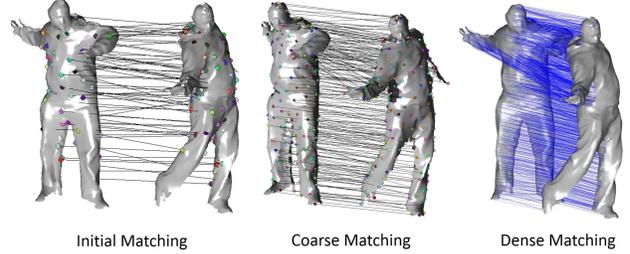


Figure 4. Matching steps.

match points on \mathcal{M}_1 and \mathcal{M}_2 that are not uniformly sampled. In addition, it is robust to non-isometric deformations.

The idea of combining ICP and graph matching is presented in the deformable graph matching paper [38]. Although we are largely inspired by [38] and combines graph matching and ICP, our algorithm is different from theirs in several aspects. First, our target is nonrigid 3D shape matching, as opposed to 2D image matching. Second, we perform ICP in high-dimensional spectral space as in [21, 23, 26], as opposed to performing in standard Euclid space. Third, our aim of fusion is to reduce complexity of QAP and to provide constraints on nonrigid transformation, as opposed to incorporate global transformation (rotation, similarity, etc.) into graph matching.

Matching steps Based on ISR, our matching technique proceeds as follows: initial matching, coarse matching and dense matching (Fig. 4). For the initial matching, we solve a graph matching problem with a standard spectral relaxation (Eq. 9), ignoring Laplacian eigen vectors. The coarse matching refines the resulting correspondences found in the initial matching and alternates Eq. 12 and Eq. 9 for 10 times. We empirically determined K as $K = 20$ at the initial matching and $K = 10$ at the coarse matching. The relatively large K at the initial matching produces globally consistent results at the cost of local accuracy. Finally, we obtain dense correspondences by performing ICP in M dimensions [23] using Laplacian eigen vectors only.

3.4. Algorithm Summary

Our algorithm is summarized as follows:

Step 1: Sampling We perform the farthest point sampling based on the diffusion distance. The number of sampling is set to $n = m = 250$ in this paper.

Step 2: Feature description We compute local depth images at the sample points. The scale factor α , which determines a view position for local depth projection, is set to 0.05, 0.1 and 0.3.

Step 3: Initial matching We solve Eq. 9 without the Laplacian eigen basis. We choose to use $K = 20$ for constructing the pairwise affinity matrix. We set the scale factor

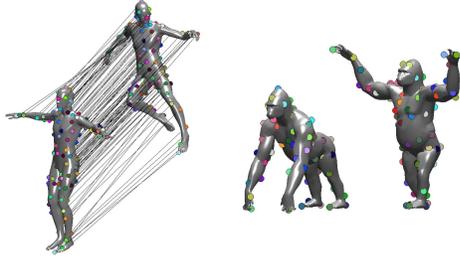


Figure 5. Coarse matching results. Randomly selected points from among 250 points are shown.

as $\sigma_p = 1$.

Step 4: Coarse matching We alternate Eq. 12 and Eq. 9 for 10 times. We choose to use $K = 8$. We set parameters to $\sigma_p = 1$ and $\sigma_e = 1$.

Step 5: Dense matching We perform ICP in M dimensions. We set $M = 100$ in this paper.

4. Results

We tested our technique on TOSCA [4], SCAPE [3], animal [30] and multi-view reconstruction [1, 16, 34] datasets. For the meshes in the multi-view reconstruction datasets, we sub-sampled them to around $20k$ vertices. The scans used in this paper typically contain 10-100 holes. We quantitatively evaluated our algorithm based on the protocol proposed in Surface Correspondence Benchmark [17] and compared the results with other state-of-the-art methods.

Wide range of models Figures 5, 8 and 9 show the ability of our method to match a wide range of models. Our matching technique can establish dense correspondences even when the models contain holes and inconsistent boundaries (Fig. 9). Furthermore, our method is robust to non-isometric deformations (Figs. 8 (a), (b) and 9) and local geometric changes such as garment motions (Fig. 8 (c)). Our technique almost can produce correct correspondences even for the shapes with different topologies (Fig. 8 (d)).

Comparisons In Fig. 6, we compare our technique with the coarse-to-fine technique (C2F) [28]. C2F produces a result with the right arm and the left arm flipped, whereas our method generates a correct result.

Figure 7 quantitatively compares our method (SAM) with Blended Intrinsic Maps (BIM) [17], Mobius Voting [19], GMDS [5] and Heat kernel map matching (HKM) [25]. We also compare the results obtained using our method but with different local descriptors: single-scale local depth maps (Single-scale), local depth maps using PCA orientations [22] (PCA orientation) and spin images [15]. Note that, in this comparison, we do not allow symmetric flips when computing geodesic errors. As can be seen in Fig. 7, our result is comparable to BIM and Mobius Voting which are robust to the presence of symmetries. Overall, BIM is slightly better than our technique. This is probably

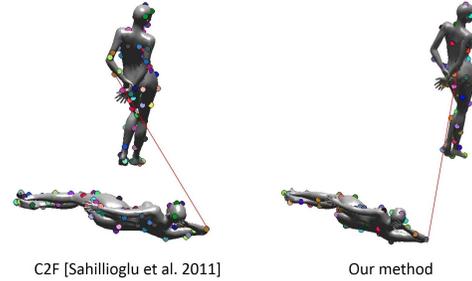


Figure 6. Comparison to C2F [28]. C2F technique produces a result with the right arm and the left arm flipped, whereas our method generates a correct result.

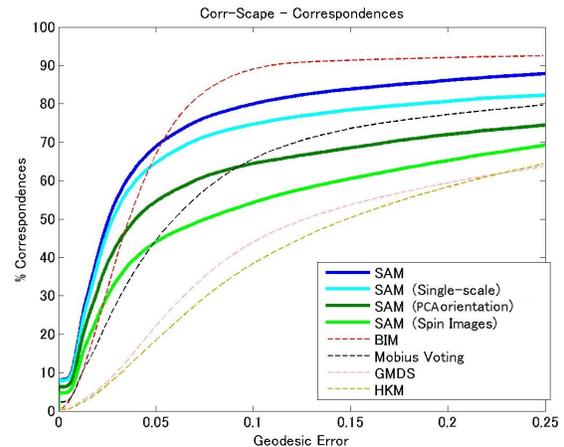


Figure 7. Quantitative comparisons. The graph shows geodesic errors on SCAPE dataset. SAM is the proposed technique (Symmetric-aware matching).

because BIM uses the global distortion measure based on conformal map. Our method sometimes picks inconsistent correspondences and produces partially wrong results. The advantage of our technique over BIM is that we can handle incomplete surfaces. We also remark that the advantages of our formulation over the sparse modeling approach [26] are that ours is symmetry-aware and robust to non-isometric deformations (see Fig. 9).

The result obtained using the ADD gradients as descriptor orientations is significantly better than the result using PCA orientation (Fig. 7 SAM vs SAM (PCA orientation)). Also, multi-scale description is shown to be effective. The result using spin images is the worst within our results, because this descriptor cannot discriminate global symmetry pairs.

Performance We implemented the prototype of our algorithm in Matlab with partially written in C/C++. On an Intel Core i7 3.4GHz 64-bit workstation, the overall algorithm takes approximately 2 min on a 50k mesh (Table 1). For 250 sample points, the number of nonzero values in \mathbf{K} was 50M and 8M for the initial matching and the coarse match-

Table 1. Timings (in seconds).

#V	#n	Sample	Feature	Init	Coarse	Dense	Total
8k	100	0.5	1.7	1.4/0.5	0.2/0.7	8.8	13.8
8k	250	1.1	4.2	9.3/5.2	1.1/10.0	9.4	40.3
50k	100	3.6	8.2	1.3/0.6	0.2/0.8	65.0	79.7
50k	250	8.8	19.5	9.2/6.5	1.1/9.8	68.0	122.9

#V – number of vertices; #n – number of sample points; Sample – timing for sampling; Feature – timing for computing globally-aligned local depth map; Init – timing for initial matching; Coarse – timing for coarse matching (The left value is the timing for constructing \mathbf{K} and the right is the time for graph matching); Dense – timing for dense matching.

ing (resp.), which was approximately $1/80$ and $1/600$ of the dense pairwise affinity matrix.

Limitations Because our technique heavily relies on local geometric information, it is difficult for our method to handle the models without surface details. Our method would fail if the surfaces lack local reflectional symmetry information or have many repetitive structures.

5. Conclusion

We presented a nonrigid shape matching method for establishing incomplete 3D surfaces in the presence of intrinsic symmetry. The key idea for matching shapes that exhibit intrinsic reflectional symmetry is to incorporate an orientation-aware local mesh descriptor that is sensitive to local reflectional symmetry. To solve QAP efficiently, we proposed the iterative spectral relaxation technique that fuses spectral embedding and spectral graph matching. This formulation improves robustness to non-isometric deformations because pairwise constraints can be defined in a scale-invariant manner.

In the future, we would like to apply our technique to other applications such as image matching, symmetry detection, etc. It would be interesting to incorporate a factorization algorithm of a global affinity matrix for further improving efficiency.

Acknowledgements

We thank Vladimir Kim for providing information about their Shape Correspondence Benchmark. We would like to thank Shinji Umeyama and Hideki Asoh for insightful discussions. This work was partially supported by JSPS KAKENHI Grant Number 25880033.

References

- [1] B. Allen, B. Curless, and Z. Popović. Articulated body deformation from range scan data. *21:612–619*, 2002.
- [2] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.*, 22(3):587–594, 2003.
- [3] 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 *NIPS*, 2004.
- [4] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer Publishing Company, Incorporated, 2008.
- [5] 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 Science*, pages 1168–1172, 2006.
- [6] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *In Proc. CVPR*, 2010.
- [7] R. E. Burkard, E. ?ela, P. M. Pardalos, and L. S. Pitsoulis. *The Quadratic Assignment Problem*, 1998.
- [8] W. Chang and M. Zwicker. Automatic Registration for Articulated Shapes. *Computer Graphics Forum (Proceedings of SGP 2008)*, 27(5):1459–1468, 2008.
- [9] T. Darom and Y. Keller. Scale-invariant features for 3-D mesh models. *IEEE Trans Image Process*, 21(5):2758–69, 2012.
- [10] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. In *Proceedings of the Symposium on Geometry Processing*, pages 1349–1356, 2008.
- [11] A. Dubrovina and R. Kimmel. Approximately Isometric Shape Correspondence by Matching pointwise Spectral Features and Global Geodesic Structures. *Advances in Adaptive Data Analysis*, 3(1-2):203–228, 2011.
- [12] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust Global Registration. In *SGP 2005: Third Eurographics Symposium on Geometry processing*, pages 197–206, 2005.
- [13] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. In *Proceedings of the Symposium on Geometry Processing*, pages 1449–1457, 2008.
- [14] V. Jain and H. Zhang. Robust 3D Shape Correspondence in the Spectral Domain. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006, SMI '06*, pages 19–, 2006.
- [15] A. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999.
- [16] N. Kasuya, R. Sagawa, R. Furukawa, and H. Kawasaki. One-shot Entire Shape Scanning by Utilizing Multiple Projector-Camera Constraints of Grid Patterns. 2013.
- [17] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4):79:1–79:12, 2011.
- [18] M. Leordeanu and M. Hebert. A Spectral Technique for Correspondence Problems Using Pairwise Constraints. In *Proceedings of the Tenth IEEE International Conference on Computer Vision - Volume 2*, pages 1482–1489, 2005.
- [19] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *ACM Trans. Graph.*, 28(3):72:1–72:12, 2009.
- [20] T. Liu, V. G. Kim, and T. Funkhouser. Finding Surface Correspondences Using Symmetry Axis Curves. *Computer Graphics Forum (Proc. Symposium on Geometry Processing)*, 2012.
- [21] D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated Shape Matching Using Laplacian Eigenfunctions and Unsupervised Point Registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [22] A. Mian, M. Bennamoun, and R. Owens. On the Repeatability and Quality of Keypoints for Local Feature-based 3D Object Retrieval from Cluttered Scenes. *Int. J. Comput. Vision*, 89(2-3):348–361, 2010.
- [23] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. J. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4):30, 2012.
- [24] M. Ovsjanikov, Q. Mérigot, V. Patraucean, and L. J. Guibas. Shape Matching via Quotient Spaces. *Comput. Graph. Forum*, 32(5):1–11, 2013.

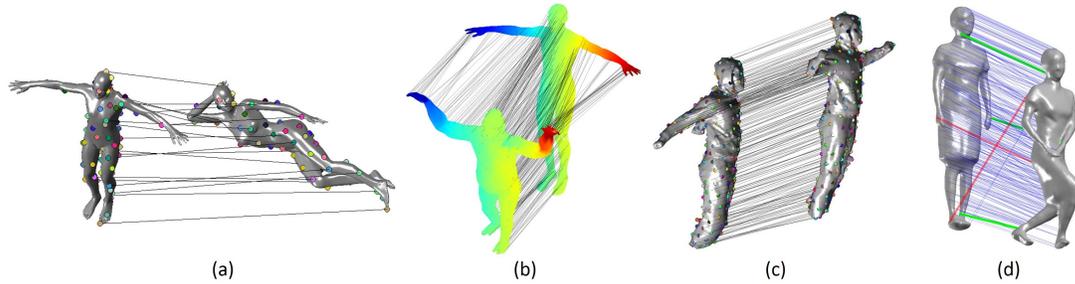


Figure 8. Robustness of our technique. Our method can match non-isometric shapes, like Male to Female (a) and Male to Gorilla (b). Our method is robust to local geometric changes such as motions of garments (c). Our technique can almost produce correct correspondences even for the shapes with different topologies (d). Green lines indicate correct matches whereas red lines indicate false matches.

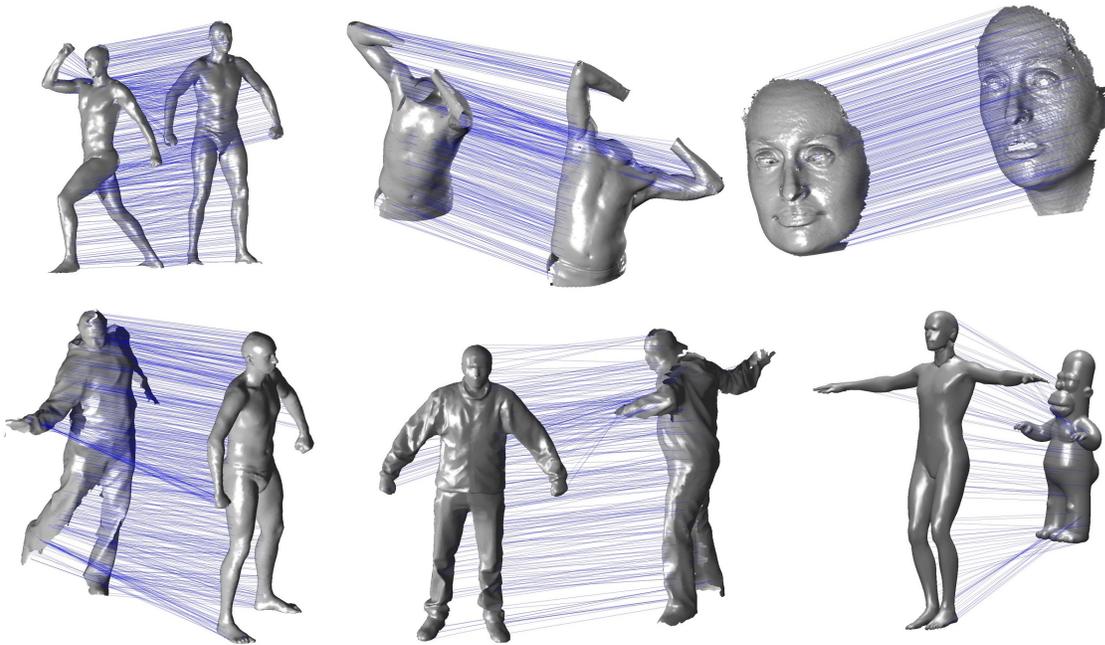


Figure 9. Dense matching results. Our technique can produce dense correspondences even for incomplete surfaces with many holes.

- [25] M. Ovsjanikov, Q. M²rigot, F. M²moli, and L. J. Guibas. One Point Isometric Matching with the Heat Kernel. *Comput. Graph. Forum*, 29(5):1555–1564, 2010.
- [26] J. Pokrass, A. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse Modeling of Intrinsic Correspondences. *Computer Graphics Forum*, 2013.
- [27] D. Raviv, M. M. Bronstein, G. Sapiro, A. M. Bronstein, and R. Kimmel. Diffusion symmetries of non-rigid shapes. In *In Proc. 3DPVT*, 2010.
- [28] Y. Sahillioglu and Y. Yemez. Coarse-to-Fine Combinatorial Matching for Dense Isometric Shape Correspondence. *Comput. Graph. Forum*, 30(5):1461–1470.
- [29] Y. Sahillioglu and Y. Yemez. 3D Shape correspondence by isometry-driven greedy optimization. In *CVPR*, pages 453–458, 2010.
- [30] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM Trans. Graph.*, 23:399–405, 2004.
- [31] J. Sun, M. Ovsjanikov, and L. J. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Comput. Graph. Forum*, 28(5):1383–1392, 2009.
- [32] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *CVPR*, pages 1185–1192, 2009.
- [33] S. Umeyama. An Eigendecomposition Approach to Weighted Graph Matching Problems. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10(5):695–703, 1988.
- [34] D. Vlastic, P. Peers, I. Baran, P. Debevec, J. Popovic, S. Rusinkiewicz, and W. Matusik. Dynamic shape capture using multi-view photometric stereo. In *In Proc. SIGGRAPH Asia*, 2009.
- [35] A. Zaharescu, E. Boyer, and R. P. Horaud. Keypoints and Local Descriptors of Scalar Functions on 2D Manifolds. *International Journal of Computer Vision*, 100(1):78–98, October 2012.
- [36] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 382–389, 2010.
- [37] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-Driven Shape Correspondence, 2008.
- [38] F. Zhou and F. De la Torre. Deformable Graph Matching. In *IEEE CVPR*, 2013.